Introduction to Cognitive Science Assignment 2: The Associative Net Due Tuesday October 27th, 2009, hardcopy, in class

How is it that you notice a mouse behind the breadbox when all that is visible is its tail? How come you notice when somone mentions your name in the midst of a buzz of noisy conversation at a Cocktail Party—whatever *that* is—even when you have no idea what they said about you? What exactly is going on when you cannot remember the name of someone you are talking to, but know that it will come to you in a minute, and it does?

These are examples of "retrieval from partial information", "recognition from noisy input", and "content addressable memory". They are all best understood in terms of massively parallel distributed processing—(M)PDP—using what are somewhat metaphorically called "Neural Networks."

Since even a two dimensional mouse represents rather a lot of information, we are going to look at recognition etc. of "one-dimensional" mice, represented by bit-vectors, or ordered sequences of 0s and 1s. (Maybe noticing a snake behind the toilet bowl would have been a better example.)

1 The Associative Net

The network you are going to work with was invented by Longuet-Higgins, Buneman, and Willshaw (see Hinton and Anderson (eds.) 1981). It is called the "Associative Net." This device illustrates three basic properties of network models which are characteristic of mechanisms involved in phenomena of human memory and attention like those mentioned above. namely:

- Non-localized storage ("Distributivity")
- Ability to recover complete stored patterns from partial or noisy input ("Graceful Degradation").
- Ability to work even in the face of damage ("Holographic Memory").

Consider the following network, connecting 16 input units each of which has a directed connection to all and only 16 output units. The connection strengths all start at 0. Input is along the top, output from the righthand side.

00000000000000000000>

Input and output patterns consist of binary vectors. So the input and output units have activities of 1 or 0. To associate an input pattern with an output pattern, we simply set to 1 all of the connections between a unit that is on in the input pattern and a unit that is on in the output pattern. If the connection was already 1, we leave it alone. Connections can never go back from 1 to 0. We will consider the "auto-associative" case where the input and the output vectors are the same—since this is the easiest case to look at the effects of noise and damage in.

Below is the net storing the pattern 0110001110101001 as input in association with itself as output:

```
0110001110101001
VVVVVVVVVVVVVVVV
0110001110101001>1
0110001110101001>1
0110001110101001>1
0110001110101001>1
0110001110101001>1
0110001110101001>1
0110001110101001>1
0110001110101001>1
```

To retrieve an association, we sum the total input, x, received by each output unit from all the on switches in its row. Then we compare x with the number y of active input units, and threshold at y. That is, if x = y, that output fires, or is turned on. If x < y, the output unit does not fire, but remains off. (If x > y there is something wrong with your program!)

Satisfy yourself that this procedure yields the same pattern 0110001110101001 as output from the memory shown above:

Input: 011000111010101 Totals: 088000888080808 Output: 011000111010101

More interestingly, if we input the "damaged" or "partial" pattern 0000001110101001, we still get 0110001110101001 as output, though with lower signal strength (6 instead of 8), and therefore a lower threshold. Check this—you should get:

Input: 0000001110101001 Totals: 0660006660606006 Output: 011000111010101 This is a bit like recognizing a mouse, snake, etc. when its head is behind the breadbox, or whatever.

Even more interestingly, if we "damage" a bit—say (7,7)—in the above memory, then from the original pattern 0110001110101001 we can still recover the complete associate, although this time we have to explicitly lower the threshold until we get something we know has been stored:

```
v
0110001110101001
0110001110101001
0110001110101001
0110001010101001<-
0110001110101001
0110001110101001
0110001110101001
0110001110101001
          v
Input:
     0110001110101001
Totals:
     0880008780808008
Output:
     0110001110101001
```

This again is somewhat like biological memory. Damage does not result in the loss of a specific memory, but to a general degradation of recall in the form of increased noise.

The homework asks you to investigate the Associative Net. As usual, you could write a program, but you can easily be the computer yourself. You can get full credit either way.

2 What you have to do

Here are some 16 bit vectors with a random 8 bits set to 1—that is, with 1s and 0s equiprobable, p = .5. Don't worry, there are far more than you need! (It is possible there are some duplicates, but its so unlikely I haven't checked.)

1101001000111001	0100010110011101	1110100100110001
0011100101110100	1000111001010011	1000001100111011
0111000111001001	0011011001100101	1100010100011101
1010001000110111	1101010011110000	1101011001101000
1011000010011011	1000110011001110	0000100101111101
1010001001101110	1010010011110100	0010011000101111
0110101000110101	1101010100100011	1110001001001101
0111000101110100	0101011010000111	1000001101101110
1000001111101010	0101011010100101	1011110010010001
1100110101100100	1010011110010100	1010011011100100
1001011110100001	1011110000011001	0011110111000010
0110010010100111	0110011010101100	0110010001111001
0111110101000001	1001100001110011	1101100010110010
1001010000110111	1110011100010100	1010011100101001
0100110110100011	0011110010100101	1101110000100110
1100110001010101	1011100100110010	0101100000110111
0011001011001110	1010000111000111	1000000101011111
1100100001111100	0001110100111010	1011100110011000
0011111000101010	1101010001101010	0011001000111110
1011101001101000	1110110001000011	0111001010110001

Customize this set by crossing out the first n vectors where n is the last digit in your student ID.

- 1. Either implement, or just draw neatly by hand, a 16 by 16 Associative net. (If you draw it you may want to make some copies of the empty net. Computers are good for doing this as well)
- 2. Load the net with the first vector in the set after the ones you have crossed out, autoassociating this vector as input with itself as output. (If you are writing a program, you should still use the above random vectors.)
- 3. Retrieve the same vector from the net, showing the total outputs on each output unit and the translation into a binary vector. (You can either write

the totals and binaries on the right of each row of the network as in the first example, or as a triple of input, totals, and output, as in later cases.)

- 4. Change the first 1 in the input vector to 0 and show that the complete pattern is still retrieved, giving totals and binary result again.
- 5. Using further vectors in sequence (you will not need many!) you are to estimate how well the net retrieves associations by iterating the following procedure.
 - (a) Add a new auto-association to the same network for the next vector on the list after the one you just added.
 - (b) Retrieve *all* the previously stored patterns including the new one from the network.
 - (c) Do this until there is some output which has at least one bit wrong.

Give all your working, showing input, total, and output, for each retrieval. State the number of patterns you managed to store up to and including the last one which caused the error.

- 6. Ideally you should repeat the whole thing using several different sets of associations with fresh empty 16x16 nets, so you get an idea of how reliable this number is. (We will use the class average as a really good estimator). You do not need to give all your working for further sets, if you do any, although you may give as much as you wish.
- 7. By making experiments along the same lines with fresh nets and sequences of vectors with higher and lower densities of 1s, try to show how this probability affects the number of associations that can be stored without more than 1 error.

Here is a bunch more vectors with four bits set, to help you, again far more than you need. Again, get rid of the first n where n is the last digit of your student ID.

101000001000100000000101000011000000000100100010000110000010000000001101100001101000000001000010100100000000011010000100000010000100110001100000000100001101000000100000110000000101000001000010100000001000001110000000000101000000100001000011000100100000000000010101000000000100100100100000000001010000000000010101000000000100100100000000100100100000000000111000000	11 10 01
01000011000001000000000011011000011010000000010000010100100000000001101000010000001000010011000110000000010000110100000010000011000000000000000000000000000000	10 01
1000001010010000000000110100001000000100001011000110000000010000110100000010000011000000000000000000000000000000	01
100011000000001000011010000000100000110001000100000100001010000000100000111000000000101010000100010010010100010010010000000000010101010000010011000010000100100100001000000111000001	
0100000100001010 0000001000001110 0000000101010 000100010010001 1000100100100000 00000010101000 000010011000010 0001001001000010 00000111000010	00
0001000010010001100010010010000000000010101010100000100110000100000100100100001000000111000001	10
0000100110000100 0001001001000010 00000111000001	00
	00
101100001000000 0010100100001000 0000001010010	01
0000100100010100 10100000100001 01000000	00
011000010000100 100000110000100 00010100001	00
100011000100000 11000000000011 0000010100	10
000010001001001 000000100010101 00100000110000	10
011010100000000 001100000110000 1101000010000	00
001000000001011 0100011000010000 00000110100001	00
0100010000001010 001110000000001 0011010010	00
100001000001100 110000011000000 1101000000	00
0000010011000100 010010000000110 0001010010	10
1000000011010000 0001000000101100 0001100100	00
0001100001000001 100100100000010 01101000100000	

(If you need vectors with other densities, just fake them up by making some appropriate number of these 0s into 1s, or 1s into 0s.)

Give your working—you can use intelligent guesswork as well as experiment for this part *if you tell us your reasoning*.

- 8. You have seen that the Willshaw net can recover the entire associated output for an input that has had one or more bits "hidden" or "damaged." Try to determine how the number of associations stored affects this ability to retrieve the complete associate from damaged input in comparison with retrieval from perfect input. Again, experiments or intelligent guesswork can be used but give your evidence in both cases.
- 9. In some sense the Associative Net "hallucinates" the whole snake when it "sees" part of it. There is no difference apart from the noise level. But it is important for a mongoose to know the difference between a snake and a snake behind the toilet-bowl, if mongooses are not to spend a lot of time bumping into things. What does this suggest to you concerning the possible psychological reality of the associative net? What is it good for if

it can't tell the difference between a snake and a snake behind the toilet? (*Hint:* Snakes are important things to notice, compared with other objects, so you may want to think in terms of a mechanism that separates *becoming aware* of snakes from the business of deciding *what to do* in a particular situation.)

Note: Write your name, school ID, and email address at the top of your homework. If you have multiple pages, *staple them together!*

Reference

Willshaw, David: (1981), "Holography, Association and Induction," in Geoffrey Hinton and James A. Anderson, (eds.), *Parallel Models of Associative Memory*, Hillsdale NJ: Erlbaum. 83-104.

mjs October 18, 2009