WORKSHEET FOR EXERCISES FROM CHAPTER 24

EXERCISE 24.1

1. **What feature is being detected by Hidden Unit 1?**

Hidden unit 1 only turns on (i.e., generates activity of 0.9 or higher) to the following patterns: 2, 3, 5, and 9. These are the only patterns in the training set that have one, and only one, input unit that is turned on. Therefore this hidden unit might be called a "1 on" detector.

2. **How do the connection weights and bias for Hidden Unit 1 permit it to detect this feature? (Remember that when an input unit is "off", it has an activation value of –1.)**

The "1 on" detector works because all of the connection weights have nearly the same value (0.58). When only one input unit is on, the signal from this unit is cancelled by the "off" signal from one of the other input units. The two remaining input units send a combined signal of about –1.16 to the hidden unit. This is nearly equal to μ, and will turn the hidden unit on. For all other patterns, this balance will not appear, and the net input will be far enough from μ so that the hidden unit does not turn on.

3. **What feature is being detected by Hidden Unit 2?**

Hidden unit 2 only turns on to patterns 8, 12, 14, and 15. These are the only patterns that have exactly 3 input units on. Therefore this hidden unit might be called a "3 on" detector.

4. **How do the connection weights and bias for Hidden Unit 2 permit it to detect this feature?**

Again, the four connection weights are nearly the same (approximately 0.62). When three inputs are turned on, one of these three signals will be cancelled by the "off" signal being sent by the fourth input unit. The other two "on" signals will work to produce a net input of about 1.24, which is nearly equal to μ which will turn the unit on. For all other kinds of input patterns, this balance will not occur, and the net input will be far enough from μ to prevent this hidden unit from turning on.

5. **How are the features being detected by the two hidden units combined in order for the output unit to make the correct response to the four-bit parity problem?**

Basically, the two hidden units described above only respond to odd-parity problems. ONe hidden unit responds to all patterns with 1 input on, while the other responds to all patterns with 3 inputs on. These define all of the possible odd-parity problems in the training set. So, if either one of the hidden units is on, the pattern must be odd parity. The connection weight from each hidden unit to the output unit has a value of 1, which is

also equal to the output unit's μ. so, when either hidden unit turns on – and only one will turn on at a time – the output unit will turn on, indicating an odd parity pattern. When an even parity problem is presented, both hidden units will be off, the net input to the output unit will be 0, and this will be far enough from μ to keep the output unit from turning on.