WORKSHEET FOR EXERCISES FROM CHAPTER 14

EXERCISE 14.1

1. **Examine the responses of the network to the training set, as well as the errors computed for each output unit and each training pattern. Your network should be having difficulties with four of the logical operations. What logical operations are causing the perceptron difficulty? (Note: if an output unit is generating three correct responses for a logical operation, but is generating an error for the fourth response, then we will say that this operation is posing a problem, simply because the output unit has failed to respond correctly to all of the input patterns.)**

There are four output units that are generating errors for at least some of the input patterns: units 8, 12, 14, and 15. These correspond to the logical operations ~(A ^ B), (A ⊃ B), (B ⊃ A), and (A ∨ B). Interestingly, these are the only four logic problems that require the network to make 3 true responses and one false response.

2. **Focus on the errors being made by the perceptron in more detail. For each output unit that is making at least one error, describe the output unit's response to each of the four input patterns.**

For the problem ~(A ^ B), the network is supposed to turn on to the first three patterns, and off to the fourth. All of the responses are in the right direction, but are not extreme enough. That is, the network only generates a response of 0.8 or so for the "on" patterns, and a response of 0.21 for the off pattern.

A similar pattern emerges for the problem (A ⊃ B). All of the responses are in the right direction. Two of these are extreme enough to be correct. The other two are errors. The network is supposed to turn on to pattern 2, but only generates activity of 0.78. The network is supposed to turn off to pattern 3, but its activity is too high: 0.37.

The pattern for (B ⊃ A) is the same as the previous, except the responses for patterns 2 and 3 are reversed. That is, the network is supposed to turn off to pattern 2, but generates activity of 0.40. The network is supposed to turn on to pattern 3, but its activity is too low: 0.82.

A similar pattern emerges for the problem (A ∨ B). All of the responses are in the right direction. One "on" response (for pattern 4) is high enough to be correct. Two others (patterns 2 and 3) are about 0.8, which is in the right direction, but not high enough to be "on". The first pattern is supposed to be "off", but the activity is 0.28, which is still too high.

3. **For each output unit that is making at least one error, examine the threshold for that unit as well as the two connection weights that are feeding into that unit. Use this information, in the context of your description of responses in question 2, to explain why the output unit is not responding correctly.**

For all of these four output patterns, the combinations of inputs and bias values for the error-producing patterns generate net inputs that are in a "grey area" that is either not far enough away from 0 to turn off completely when desired, or not close enough to 0 to turn on completely when desired. One of the problems seems to be, geometrically speaking, that when three units need to be turned on, the "band" carved by the Gaussian is too narrow to fit them all in close enough together for all of them to turn on.

4. **On the basis of your answer to question 3, is it possible in principle for these incorrect output units to eventually learn to respond correctly, or are they doomed to eternal failure? If you think that they can respond correctly, then explain why, and return to the program to try to validate this belief empirically. (Hint: I don't recommend this latter approach!)**

While the responses are in the right direction, the network is not likely to generate the correct responses for all of the patterns in the troublesome logic problems. Again, this is because the carving of the pattern space made by the value unit seems to have trouble dealing with patterns in which there are three true responses, which is the case for all four of these logic problems. The triangle of "on" patterns is too large to fit between the two cuts made in the pattern space by the value unit.

5. **Compare the difficulties that you found for this network with the difficulties that you found for the network in Chapter 12. Are they the same? If not, then why not? What are the implications of this comparison for choosing a perceptron to solve a particular pattern classification problem?**

WORKSHEET FOR EXERCISES FROM CHAPTER 14

EXERCISE 14.1

6. **Examine the responses of the network to the training set, as well as the errors computed for each output unit and each training pattern. Your network should be having difficulties with four of the logical operations. What logical operations are causing the perceptron difficulty? (Note: if an output unit is generating three correct responses for a logical operation, but is generating an error for the fourth response, then we will say that this operation is posing a problem, simply because the output unit has failed to respond correctly to all of the input patterns.)**

There are four output units that are generating errors for at least some of the input patterns: units 8, 12, 14, and 15. These correspond to the logical operations ~(A ^ B), (A ⊃ B), (B ⊃ A), and (A ∨ B). Interestingly, these are the only four logic problems that require the network to make 3 true responses and one false response.

7. **Focus on the errors being made by the perceptron in more detail. For each output unit that is making at least one error, describe the output unit's response to each of the four input patterns.**

For the problem ~(A ^ B), the network is supposed to turn on to the first three patterns, and off to the fourth. All of the responses are in the right direction, but are not extreme enough. That is, the network only generates a response of 0.8 or so for the "on" patterns, and a response of 0.21 for the off pattern.

A similar pattern emerges for the problem (A ⊃ B). All of the responses are in the right direction. Two of these are extreme enough to be correct. The other two are errors. The network is supposed to turn on to pattern 2, but only generates activity of 0.78. The network is supposed to turn off to pattern 3, but its activity is too high: 0.37.

The pattern for (B ⊃ A) is the same as the previous, except the responses for patterns 2 and 3 are reversed. That is, the network is supposed to turn off to pattern 2, but generates activity of 0.40. The network is supposed to turn on to pattern 3, but its activity is too low: 0.82.

A similar pattern emerges for the problem (A ∨ B). All of the responses are in the right direction. One "on" response (for pattern 4) is high enough to be correct. Two others (patterns 2 and 3) are about 0.8, which is in the right direction, but not high enough to be "on". The first pattern is supposed to be "off", but the activity is 0.28, which is still too high.

8. **For each output unit that is making at least one error, examine the threshold for that unit as well as the two connection weights that are feeding into that unit. Use this information, in the context of your description of responses in question 2, to explain why the output unit is not responding correctly.**

For all of these four output patterns, the combinations of inputs and bias values for the error-producing patterns generate net inputs that are in a “grey area” that is either not far enough away from 0 to turn off completely when desired, or not close enough to 0 to turn on completely when desired. One of the problems seems to be, geometrically speaking, that when three units need to be turned on, the “band” carved by the Gaussian is too narrow to fit them all in close enough together for all of them to turn on.

9. **On the basis of your answer to question 3, is it possible in principle for these incorrect output units to eventually learn to respond correctly, or are they doomed to eternal failure? If you think that they can respond correctly, then explain why, and return to the program to try to validate this belief empirically. (Hint: I don’t recommend this latter approach!)**

While the responses are in the right direction, the network is not likely to generate the correct responses for all of the patterns in the troublesome logic problems. Again, this is because the carving of the pattern space made by the value unit seems to have trouble dealing with patterns in which there are three true responses, which is the case for all four of these logic problems. The triangle of “on” patterns is too large to fit between the two cuts made in the pattern space by the value unit.

10. **Compare the difficulties that you found for this network with the difficulties that you found for the network in Chapter 12. Are they the same? If not, then why not? What are the implications of this comparison for choosing a perceptron to solve a particular pattern classification problem?**