CPSC 633 Exam #1 (3/07/2011, Tue)[1]

Last name: ________________, First name: ________________

Time: **9:35am–10:50am (75 minutes)**, Total Points: **100**

Subject	Score
Supervised Learning	/20
Decision Tree Learning	/20
Neural Networks	/30
Reinforcement Learning	/30
Total	/100

- Leave one empty seat between you and your next person.
- Show your student ID on your desk.
- Sit according to your handedness.
- Be as **succinct** as possible. If you write too much stuff and if you make both correct statements and incorrect statements, **you will get 0 point**.
- Do not repeat the question. Just say "It is because ...".
- Read the questions carefully to see what kind of answer is expected (*explain blah* **in terms of ...** *blah*).
- If you feel that the question is not specific enough, please ask.
- Solve all problems.
- This is a closed-book, **open-note** (your hand-written cheat sheet) exam.

[1] Instructor: Yoonsuck Choe.

1 Supervised Learning

Question 1 (10 pts): (1) Is it possible for an input $x \in \mathcal{D}$ to be incorrectly classified by a hypothesis h in $VS_{H,D}$, where $\mathcal{D}$ is the full input space, $VS_{H,D}$ is the version space, H is the hypothesis space and D is the training set ($D \subset \mathcal{D}$)? That is, can an x exist in $\mathcal{D}$ so that $h(x) \neq c(x)$, where $h \in VS_{H,D}$ and c is the true concept (or target function)? (2) Explain why.

Question 2 (10 pts): (1) Is having a hypothesis space that has a high VC dimension good, bad, or neither? (2) Explain in terms of both (a) richness or expressive power of the hypotheses, and (b) sample complexity. (Hint: first start by asking if higher VC dimension leads to higher/lower richness and higher/lower sample complexity.)

2 Decision Tree Learning

Question 3 (10 pts): (1) Is the information gain measure related to the inductive bias in decision tree learning algorithm ID3? (2) Explain why or why not. (Hint: Start by thinking about what is the inductive bias in ID3.)

Question 4 (10 pts): In ID3, uncertainty typically decreases when a particular attribute is checked. (1) Name the quantity that represents "the amount of decrease in uncertainty". (2) Describe a scenario where this quantity can be 0.

3 Neural Networks

Question 5 (10 pts): (1) Can you apply gradient descent learning when the activation function is a step function instead of a sigmoid function? (2) Explain why or why not.

$$\mathrm{step}(x) = \begin{cases} 1 & \text{if } x \geq 0 \\ 0 & \text{if } x < 0 \end{cases}$$

Question 6 (10 pts): (1) Is the following perceptron with a single input x and a bias 1 able to learn to correctly classify the input set shown below ("o" and "x" are the two different classes)? (2) Explain why or why not. (3) For this perceptron, what is the dimensionality of the separating hyperplane? (Hint: Write down the perceptron formula in full to see what form the separating hyperplane takes. Also, figure out the input dimensionality.)

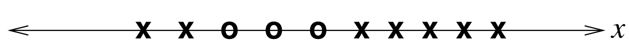

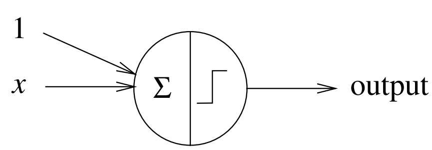

Question 7 (10 pts): Why can online training (as opposed to batch training) help backpropagation learning escape from local minima? Explain in terms of the error surface.

4 Reinforcement Learning

Question 8 (10 pts): The initial definition of $Q(s, a)$ is

$$Q(s, a) \equiv r(s, a) + \gamma V^*(\delta(s, a)).$$

There are several terms that are unknown to the learning system: $r(s, a)$, $V^*(s')$, and $\delta(s, a)$. Explain how Q-learning overcomes these unknowns (Hint: compare to $\hat{Q}(s, a)$).

Question 9 (10 pts): Explain each step of the following derivation, (1) from Step 0 to Step 1, and (2) from Step 1 to Step 2:

$$\begin{aligned} Q(s,a) &\equiv E[r(s,a) + \gamma V^*(\delta(s,a))] \qquad \{\text{Step 0}\} \\ &= E[r(s,a)] + \gamma E[V^*(\delta(s,a))] \qquad \{\text{Step 1}\} \\ &= E[r(s,a)] + \gamma \sum_{s'} P(s'|s,a) V^*(s') \qquad \{\text{Step 2}\} \end{aligned}$$

Question 10 (10 pts): What is the main difference between Q-learning and SARSA?