

# Resit Exam Machine Learning

## 2010-2011

10 June 2011

This exam is **open book**: you can use Tom Mitchell's "Machine Learning" as well as the lecture slides and any notes you've taken. You can use a calculator.

Answers are allowed in Dutch and English.

**Good luck!**

### Questions

#### 1. Short answers (15 Points)

- (a) (**True or False**): The back-propagation algorithm learns a globally optimal neural network with hidden layers.
- (b) (**True or False**): The error of a hypothesis measured over its training set provides a pessimistically biased estimate of the true error of the hypothesis (i.e., the real-world performance of the hypothesis would be better than the training set performance).
- (c) (**True or False**): Overfitting is more likely when the hypothesis space is small
- (d) (**True or False**): Because gradient descent is guaranteed to find a local optimum, it cannot overfit.
- (e) (**True or False**): When training neural networks, the network weights are updated more often in incremental or stochastic gradient descent than in batch gradient descent.

#### 2. Decision Trees (20 Points)

Summertime is festival time: Pinkpop, Pukkelpop, Parkpop, to name but a few. To help you choose which of the acts to go see at the next festival you're visiting, you will develop a model based on the data in table 1.

- (a) What are the entropies  $H(\text{Like}|\text{Style} = \text{Britpop})$  and  $H(\text{Like}|\text{Set length} = 30)$ ?<sup>1</sup>

---

<sup>1</sup>Note: if your calculator can't do  $\log_2$ , use one of the following equations:  $\log_2(x) = 1.44 \cdot \ln(x)$  or  $\log_2(x) = 3.32 \cdot \log_{10}(x)$ . If you didn't bring a calculator, you may give the answer as an expression.

Year	Age	Sex	Length (mm)	Weight (g)
1996	1	Male	30	0.5
1996	1	Female	30	0.5
1996	2	Male	35	0.8
1996	2	Female	35	0.8
1996	3	Male	40	1.2
1996	3	Female	40	1.2
1996	4	Male	45	1.8
1996	4	Female	45	1.8

- (A) What rationale would this decision involve taking appropriate notice of the fact that the defendant is a minor?
- (B) Does the law decision now cited would be based on the fact that the defendant is a minor?
- (C) What is the rationale for this decision? What is the meaning of the word "appear" as the percentage of records that would be submitted?

2. Repeat the questions (10-15).

- (A) Consider whether this decision is based on the extent of the defendant's knowledge of the fact that the defendant is a minor. Is the defendant's knowledge of the fact that the defendant is a minor relevant to the decision?
- (B) Does the law decision now cited would be based on the fact that the defendant is a minor?
- (C) What is the rationale for this decision? What is the meaning of the word "appear" as the percentage of records that would be submitted?

3. Repeat the questions (10-15).

- (A) Consider whether this decision is based on the extent of the defendant's knowledge of the fact that the defendant is a minor. Is the defendant's knowledge of the fact that the defendant is a minor relevant to the decision?
- (B) Does the law decision now cited would be based on the fact that the defendant is a minor?
- (C) What is the rationale for this decision? What is the meaning of the word "appear" as the percentage of records that would be submitted?

## 5. Instance-Based Learning (15 Points)

- (a) The following data set (Fig. 1) shows classification data with two classes: plus and minus. Which of the two nearest neighbours classifiers has the largest Leave-One-Out Cross-validation error on this data set?
- 1-NN
  - 3-NN




Figure 1: A classification data set

- (b) Suppose that you want to use k-nearest neighbour on a data set that contains a non-numeric attribute 'weight' which can be *light*, *medium* or *heavy*. Could you re-code this attribute to be able to use it? How?
- (c) The same for an attribute 'material' which can have the values *aluminium* or *plastic*.

## 6. Hypothesis Comparison and Cross Validation (15 Points)

- (a) On a much larger data set than that in table 1, someone has developed a classifier that correctly classifies 700 out of 1000 test examples. Give an estimate of the true error of this classifier, and a 90% confidence interval around that estimate (you may give your confidence interval in the form of an expression).<sup>2</sup>
- (b) In the above problem, if you increase the size of the test set by a factor of 10, assuming a similar error rate of 30%, what happens to the confidence interval?
- (c) When tuning the parameters of a learning algorithm (e.g., the number of neighbours in kNN) to a given application, why is it necessary to use cross-validation or a test set?

<sup>2</sup>The table of z-values:

N%:	50%	68%	80%	90%	95%	98%	99%
$z_N$ :	0.67	1.00	1.28	1.64	1.96	2.33	2.58