

6.034 Quiz 4

3 December 2014

Name	
Email	

Circle your TA (**for 1 extra credit point**), so that we can more easily enter your score in our records and return your quiz to you promptly.

Josh Blum

Elisa Castañer

Pedro Cattori

Martin Gale

Malcom Gilbert

Dylan Holmes

Jessica Noss

Duncan Townsend

Siyao Xu

Problem number	Maximum	Score	Grader
1	50		
2	50		
Total	100		

SRN	6		
-----	---	--	--

There are 9 pages in this quiz, including this one. As always, open book, open notes, open just about everything, including a calculator, but no computers.

Problem 1: Adaboost (50 points)

Part A: Two rounds of boosting (32 points)

A1 (22 points)

You have six training points (A, B, C, D, E, F) and five classifiers (h_1, h_2, h_3, h_4, h_5) which make the following misclassifications:

Classifier	Misclassified training points (A, B, C, D, E, F)					
h_1	A			D		F
h_2				D		
h_3		B	C			
h_4	A	B				F
h_5		B	C	D		

Perform two rounds of boosting with these classifiers and training data. In each round, pick the classifier with the **lowest error rate**. Break ties by picking the classifier that comes first in this list: h_1, h_2, h_3, h_4, h_5 . Space for scratch work is provided on the following page.

	Round 1	Round 2
weight _A	1/6	
weight _B		
weight _C		
weight _D		
weight _E		
weight _F		
Error rate of h_1		
Error rate of h_2		
Error rate of h_3		
Error rate of h_4		
Error rate of h_5		
weak classifier (h)		
classifier error (ϵ)		
voting power (α)		

Space provided for scratch work

A2 (6 points)

Three of the training points (B, D, F) have been selected below. For each one, decide whether the ensemble classifier $H(\vec{x})$ produced after two rounds of boosting misclassifies or correctly classifies that point. Circle the best answer in each case. If the answer can't be determined from the available information, circle "Can't tell" instead.

Training point	Classification by ensemble classifier		
B	Correctly classified	Misclassified	Can't tell
D	Correctly classified	Misclassified	Can't tell
F	Correctly classified	Misclassified	Can't tell

A3 (4 points)

Suppose you continue the Adaboost procedure from Part A for a **total** of 2014 rounds. (You may assume it doesn't terminate before then.) If you always pick the classifier with the **lowest** error rate, which training data point will have the **smallest** weight at the end of the 2014th round? (Circle one.)

A B C D E F

Part B (18 points)

This section consists of questions about Adaboost *in general*—they do not rely on the preceding section. Decide whether each of the statements below is true or false. Circle the one best answer in each case.

- True False 1. Adaboost accounts for outliers by lowering the weights of training points that are repeatedly misclassified.
- True False 2. When you update weights, the training point with the smallest weight in the previous round will always increase in weight.
- True False 3. Four weak classifiers that make disjoint errors can always be used to construct a perfect classifier by picking $\alpha_1 = \alpha_2 = \alpha_3 = \alpha_4 = 1$.
- True False 4. In each round, any weak classifier with an error rate less than $\frac{1}{2}$ must misclassify fewer than half of the training points.
- True False 5. After two rounds of boosting, the ensemble classifier $H(\vec{x})$ always behaves like the conjunction (AND) of two weak classifiers.
- True False 6. If you have 100 training points, the updated weight of a training point can never be greater than $\frac{3}{4}$.
- True False 7. A weak classifier will never have an error rate greater than $\frac{1}{2}$.
- True False 8. Once a weak classifier is picked in a particular round, it will never be chosen in any subsequent round.
- True False 9. Weak classifiers with error rates (ϵ) close to 100% are assigned negative voting powers (α).

Problem 2: Bayesian Inference (50 points)

★ Note: This problem consists of four independent subproblems. ★

Part A: Bayes rule and model selection (10 points)

Suppose you have three coins in a bag: the first coin is fair, with HEADS on one side and TAILS on the other. The second has HEADS on both sides. The third has TAILS on both sides.

Your friend takes a coin out of the bag at random and flips it. When it lands, you observe that the side of the coin facing up is HEADS. What is the probability that this is the fair coin?

$$\frac{2}{3}$$

(A)

$$\frac{1}{2}$$

(B)

$$\frac{1}{3}$$

(C)

$$\frac{1}{4}$$

(D)

Check your intuition: You should treat each coin as a different model M_1, M_2, M_3 for explaining the evidence E of having HEADS on one side. Thus, the probability that the coin is fair is $P(M_1|E)$, which you can compute using Bayes' Law.

Space for scratch work.

Part B: Sensitivity and specificity (20 points)

Senioritis is a rare but treatable condition that occurs in about 20 out of every 10,000 individuals in the general population. A recently developed test for senioritis—the PUNT scan—is 99.5% sensitive and 80% specific. (This means that 995 out of 1000 people with senioritis correctly test positive, and 800 out of every 1000 people without senioritis correctly test negative. The PUNT scan always reports either “positive” or “negative”.)

For notation, we can let D be the variable “You have senioritis” and let T be the variable “You test positive for senioritis”. Then the information above is:

$$\begin{aligned}P(D) & \quad 20 \text{ out of } 10,000 \\P(T \mid D) & \quad 995 \text{ out of } 1,000 \\P(T \mid \bar{D}) & \quad 800 \text{ out of } 1,000\end{aligned}$$

Part B1 (10 points) What is the probability of obtaining a positive test result, regardless of whether you have senioritis?

The marginal probability of a positive test result is *approximately* (circle one):

0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100%

You must show your work to receive credit. Write down the equations you intend to solve, if any, and indicate what values you're plugging in. You probably won't need a calculator, since you only need an approximate final answer.

Part B2 (10 points) Suppose your PUNT scan returns a positive result. In this case, the probability that you indeed have senioritis is *most nearly* (circle one):

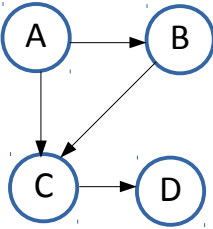
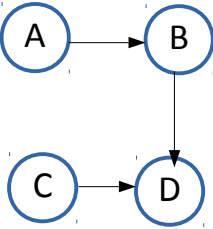
0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100%

You must show your work to receive credit. Write down the equations you intend to solve, if any, and indicate what values you're plugging in. You probably won't need a calculator, since you only need an approximate final answer.

Part C: Applying d-separation (16 points)

In the figure below, there are two Bayes nets and some independence statements. For each of the statements below and each Bayes net, circle TRUE if the statement is true for the net, and FALSE if the statement is false for the net.


Note: Assume that the only independence statements that are true are the ones enforced by the shape of the network.

		
A is independent of C.	TRUE FALSE	TRUE FALSE
Given C, A is independent of D.	TRUE FALSE	TRUE FALSE
$P(B DAC) = P(B AC)$	TRUE FALSE	TRUE FALSE
Assuming all of the variables are boolean, how many parameters does each Bayes net have? (The number of parameters is the total number of entries in all probability tables.)		
# of parameters		

Part D: Make no assumptions (4 points)

In exactly **one** of the Bayes nets below, **none of the variables are independent**. Circle it.

Hint: A Bayes net that makes no independence assumptions is a model that is consistent with every joint probability table. You may find it useful to consider how many parameters it must have as a result.



Problem 3: Spiritual and Right-Now

Circle the **one best** answer for each of the following questions. There is **no penalty for wrong answers**, so it pays to guess in the absence of knowledge.

1 Boris Katz's START system was an inspiration leading to:

1. The Deep Blue chess-playing program.
2. The introduction of subliminal advertising.
3. The Heritage system for genetic counseling.
4. The Google and Stanford systems for describing actions in pictures.
5. The Siri and Watson systems for question answering.

2 Probabilistic programming has been used with success to:

1. Analyze tweets.
2. Analyze voting data.
3. Analyze satellite data.
4. Improve athletic training.
5. Improve use of PowerPoint presentations in education.

3 A key idea behind the Genesis story-understanding system is the use of:

1. Rules and search for reasoning and conceptual analysis.
2. Constraint propagation to anticipate story outcomes.
3. Support vector machines to find common features in story collections.
4. Boosting to determine whether stories are interesting.
5. A small collection of standard descriptive frames to summarize stories.

4 A key idea behind the success of deep neural nets is:

1. Unsupervised, layer by layer training, aimed at producing outputs that are the same as inputs.
2. Use of dot-product kernel functions to transform difficult problems into easier problems.
3. Coupling small learning rates to massive computing to avoid local maxima.
4. Linear dependence of the net depth on the number of classes to be recognized.
5. Use of back propagation, accelerated by using a larger rate multiplier in layers closer to inputs.

5 Developing a most probable story-generation model from a collection of stories involves:

1. Recognizing similar stories using a nearest neighbor classifier.
2. Recognizing similar stories using a boosting classifier.
3. Assuming all possible story models are equally likely.
4. Averaging over a large number of stories using crowd sourcing to decide what to merge.
5. Using a carefully tuned prior that is a function of model size to guide Bayesian structure search.

6 When a rat stop moving on a track, activity in the hippocampus:

1. Stops completely.
2. Depends on whether the rat is hungry or recently fed.
3. Can be decoded to determine whether the rat is happy or stressed.
4. Sometimes replays activity as if the rat is moving forward, sometimes as if moving backward.
5. Bounces around, suggesting the rat has random thoughts about where it is and what it is doing.