6.034 Quiz 4

5 December 2018

Name	
Email	

For 1 extra credit point: Circle the TA whose recitations you attend so that we can more easily enter your score in our records and return your quiz to you promptly.

Suri Bandler   **Marie Feng**   **Kifle Woldu**

Sanchit Bhattacharjee   **Ariel Jacobs**   **Matt Wu**

Alex Charidis   **Victoria Longe**   **Richard Yip**

Samir Dutta   **Smriti Pramanick**

Problem number	Maximum	Score	Grader
1 - Adaboost	56		
2 – Bayes Nets	44		
Total	100		

SRN	6		

There are 12 pages in this quiz, including this one, but not including tear-off sheets. Tear-off sheets with duplicate drawings are located after the final page of the quiz.

As always, open book, open notes, open just about everything, including a calculator, but no computers.

This page is intentionally blank.

Problem 1: Adaboost (56 points)

Part A: Six of Crows (40 points)

Six (6) students from the bestselling book *Six of Crows*—**R**achel (R), H**O**lden (O), **M**ahgi (M), **T**asha (T), **K**yle (K), and **S**car (S)—are headed back to the Ketterdam University District after an evening out.

The students want to use Adaboost to determine if they will each arrive home safely, and they poll patrons at five (5) establishments about how safe each of their trips home would be. For each student, the patrons at each establishment act as a weak classifier and predict the student's trip home as safe (yes) or not safe (no). The resulting **five (5) weak classifiers**—Kaelish Prince, **KP(x)**; Emerald Palace, **EP(x)**; Crow Club, **CC(x)**; Blue Paradise, **BP(x)**; The Sweet Shop, **SS(x)**—make the following misclassifications:

Weak Classifiers	Misclassified Training Points
KP(x)	R, O, M
EP(x)	R, O
CC(x)	M
BP(x)	T, K, S
SS(x)	O, M, T

For your convenience, the table is provided on a tear-off sheet after the last page of the quiz.

A1 (30 points) On the next page, complete the first three (3) rounds of Adaboost,

A1 (30 points) Fill out the table below to complete the first three (3) rounds of Adaboost, choosing the classifier with the **error rate furthest from ½**. Some of the values have been filled in for you.

Break ties according to which classifier comes first in this list: KP(x), EP(x), CC(x), BP(x), SS(x).

You can show your work for partial credit on the next page.

	Round 1	Round 2	Round 3
Weight of R (w_R)	$^1/_6$		
Weight of O (w_O)			
Weight of M (w_M)			
Weight of T (w_T)			
Weight of K (w_K)			
Weight of S (w_S)			
Error rate of KP(x) (ε_{KP})			$^{13}/_{16}$
Error rate of EP(x) (ε_{EP})			$^1/_2$
Error rate of CC(x) (ε_{CC})			$^5/_{16}$
Error rate of BP(x) (ε_{BP})			$^3/_{16}$
Error rate of SS(x) (ε_{SS})			$^5/_8$
Weak classifier chosen (h)			
Weak classifier error (ε)			
Voting power (α)			

The following rules for manipulating logarithms may be helpful.

$$\log(x)+\log(y)=\log(x*y) \qquad \log\frac{1}{x}=-\log(x) \qquad \log(1)=0$$

4

For partial credit on A1, you can show your work here.

A2 (3 points) What ensemble classifier **H(x)** did you generate after three (3) rounds of Adaboost?

H(x) =

A3 (1 points) Which weak classifier **h** contributes the most to H(x)? <u>Circle one</u>.

KP(x) **EP(x)** **CC(x)** **BP(x)** **SS(x)**

A4 (6 points) How does the H(x) produced after 3 rounds of Adaboost classify each of the following training points? <u>Circle one</u>.

O: Correctly Classified Misclassified Can't Tell

M: Correctly Classified Misclassified Can't Tell

T: Correctly Classified Misclassified Can't Tell

Part B: Ensembles? (16 points)

The students disagree about ensemble classifiers. **Circle** whether each student's statement below is correct, incorrect, or can't be determined from available information. **In one sentence explain your choice** in the box provided. Assume all Adaboost parameters, including weak classifiers, are the same as in Part A. Also assume as in Part A that the best classifier is the one with error rate furthest from ½.

B1 (4 points) Mahgi: If you performed Adaboost with no limit on the number of rounds and only used two (2) of our five (5) classifiers, then the accuracy will never exceed 80%.

Correct	**Incorrect**	**Can't Tell**

Explanation:

B2 (4 points) Rachel: Some of our five (5) classifiers will never be picked by the Adaboost algorithm.

Correct	**Incorrect**	**Can't Tell**

Explanation:

B3 (4 points) Tasha: By choosing appropriate voting powers by hand, you can make a perfect ensemble classifier out of three (3) of our five (5) weak classifiers.

Correct	**Incorrect**	**Can't Tell**

Explanation:

B4 (4 points) Kyle decides to make his own ensemble classifier from three (3) of the weak classifiers:
$$H(x) = 4*CC(x) + 2*EP(x) + BP(x)$$

Kyle: My classifier $H(x)$ is made from three (3) weak classifiers that make disjoint errors. Therefore, it has perfect accuracy on the original training data.

Correct	**Incorrect**	**Can't Tell**

Explanation:

Problem 2: Bayes Nets (44 points)

Part A: Genie in a Bayes Net (28 points)

Christina Aguilera asks 6.034 to help her write next summer's big hit song to earn her nineteenth GRAMMY nomination. She uses her understanding of the music scene to compile a probability table with six (6) boolean variables, but she needs your help to analyze the table.

A1 (4 points) Without making any assumptions about independence, how many parameters do we need in the joint probability table? (You don't need to simplify your answer.)

> *Number of parameters =*

A2 (6 points) Christina thinks that with no independence assumptions, the joint probability table is much too big. To make her model computationally feasible with her six (6) boolean variables, she defines the Bayes net below so that she can rely on the Bayes net independence assumption.

Assume the only independence statements that are true are the ones enforced by the shape of the net. How many parameters are needed for this Bayes net?

> *Number of parameters =*

For partial credit, next to each variable in the Bayes net, clearly write the number of parameters in that particular variable's probability table.
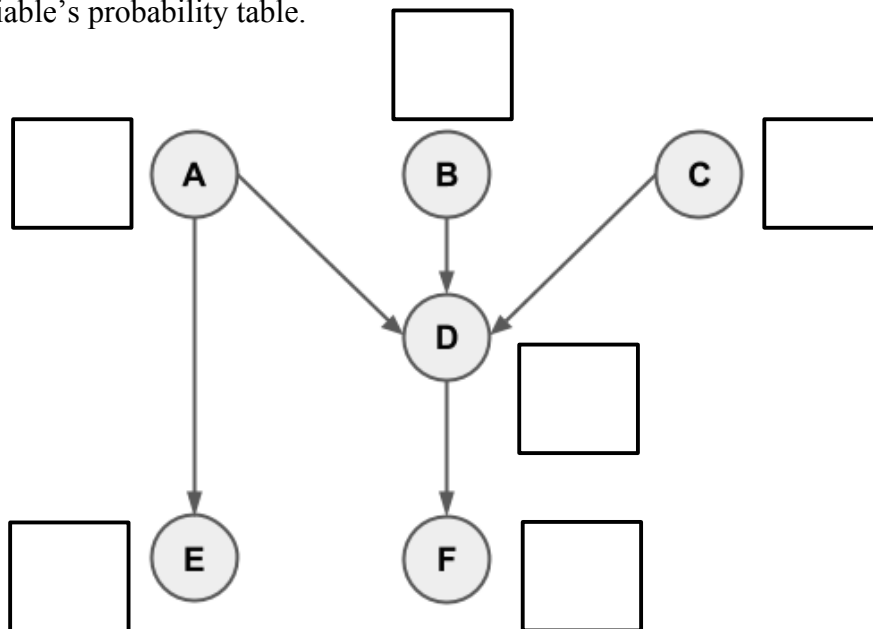
For problems A3 through A5, refer to the Bayes net on the previous page. The net is also duplicated on a tear-off sheet. Again assume the only independence statements that are true are the ones enforced by the shape of the net (all other independence statements are false).

A3 (6 points) $P(A \mid C\,E) = P(A \mid E)$. <u>Circle one</u>.

True	**False**

For partial credit on A3, show your d-separation work here.

A4 (6 points) A and D are marginally independent. <u>Circle one</u>.

True	**False**

For partial credit on A4, show your d-separation work here.

A5 (6 points) E and B are conditionally independent, given D. <u>Circle one</u>.

|	
True	**False**

For partial credit on A5, show your d-separation work here.

Part B: Hit Me Bayesian One More Time (16 points)

Jealous of Christina, Britney Spears asks 6.034 to help her write next summer's even bigger hit song so she can earn her ninth GRAMMY nomination.

She believes that whether a song is nominated for GRAMMY consideration (C) depends on whether it uses AIToTune (A) and whether it is on the Billboard Hot 100 (B). She also believes that being on the Billboard Hot 100 (B) influences whether a song's rhythm changes with a beat drop (D).

For your convenience, a copy of this Bayes net is provided on a tear-off sheet.
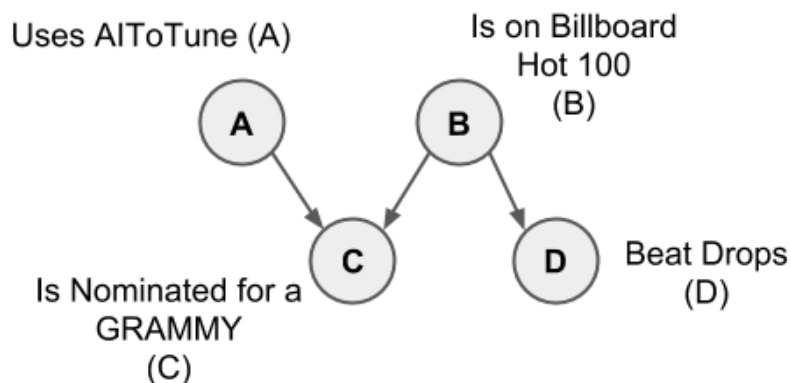
Britney wants to know her chances of being considered for a GRAMMY (C = TRUE), given that A = FALSE, B = FALSE, and D = FALSE.

In other words, Britney is computing an expression for the joint probability $P(\overline{A}\,\overline{B}\,C\,\overline{D})$. You plan to use the Chain Rule to expand the joint probability as a product of conditional probabilities with increasing numbers of given variables.

B1 (4 points) Expand $P(\overline{A}\,\overline{B}\,C\,\overline{D})$ as a product of conditional probabilities. (*Hint: Consider expanding the expression in a particular order to best take advantage of the structure of the network, but don't apply any independence assumptions to simplify your conditional probabilities yet.*)
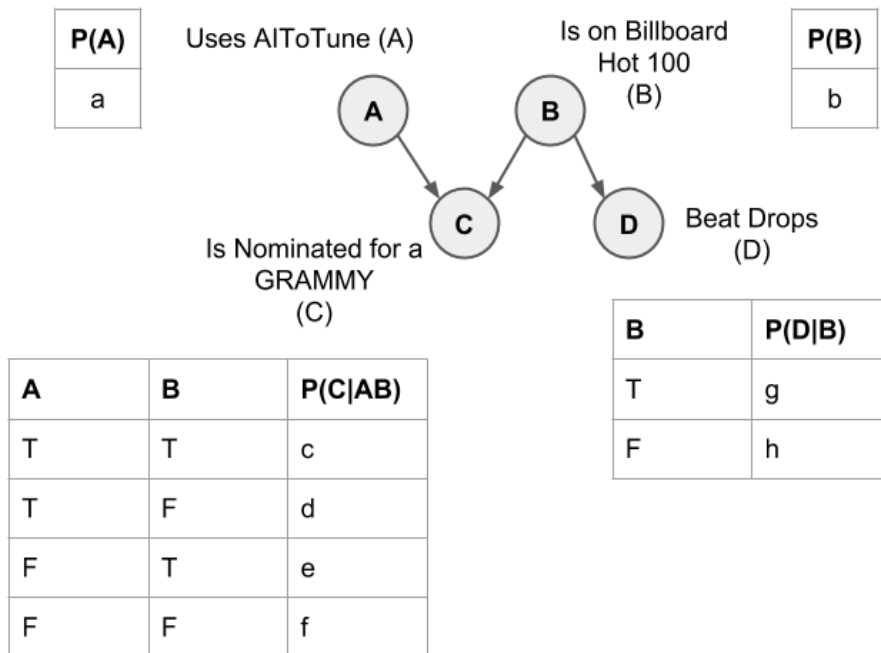
$P(\overline{A}\,\overline{B}\,C\,\overline{D})$ =

B2 (6 points) Simplify your expression in **B1** using the Bayes net independence assumption. *Your answer should be expressed in terms of conditional and marginal probabilities that can be looked up in the Bayes net.*

$P(\bar{A}\,\bar{B}\,C\,\bar{D})$ =

B3 (6 points) Brittney provides you with the parameters of the Bayes net. In the associated probability tables, each probability is represented by a lowercase letter.

For your convenience, a copy of this Bayes net with the tables is provided on a tear-off sheet.

P(A)
a

Uses AIToTune (A)

Is on Billboard Hot 100 (B)

P(B)
b

A   **B**

C   **D**

Is Nominated for a GRAMMY (C)

Beat Drops (D)

A	B	P(C|AB)
T	T	c
T	F	d
F	T	e
F	F	f

B	P(D|B)
T	g
F	h

Evaluate your expression in **B2** and write it as an algebraic expression of the parameters a through h specified in the Bayes net.

$P(\bar{A}\,\bar{B}\,C\,\bar{D})$ =

10

Problem 3: Spiritual and Right Now (6 points)

For each question, **in the box provided, write a letter** corresponding to **one** best answer **and circle the answer**. There is **no penalty for wrong answers**, so it pays to guess in the absence of knowledge.

1. Winston argued that inheritance is crucial in role frame representations because inheritance:
   a. makes it easier to put the representation into code and translate between coding languages.
   b. demonstrates how climbing the hierarchy contributes increasingly large jumps in knowledge.
   c. gives us names over the concepts and their relationships to one another.
   d. can be modeled as linking cause and effect, growing our semantic net.

2. Winston uses the term "parasitic semantics" to refer to:
   a. semantic theories riddled with conceptual holes, as if infected by parasites.
   b. supposing that a program has humanlike knowledge about labels in its vocabulary.
   c. a Marxist theory of knowledge representation developed to condemn social parasites.
   d. informal theories that gain power by mathematical mappings to Aristotelian logic.

3. Katz argued that a good way to collect information needed for constructing a large set of natural language questions is to:
   a. download relational databases from public resources, e.g., The Census or Google Scholar.
   b. gather semi-structured data from the information boxes on the top of Wikipedia pages.
   c. parse Wikipedia articles to understand what the articles are about.
   d. download encyclopedias such as World Book and Encyclopedia Brittanica.

4. Berwick argued that Aristotle's model of language, with components sound and meaning, is incomplete and requires a third component, called merge, which he referred to as a:
   a. language CPU, which processes speech into meaning and vice versa.
   b. language cache, which stores utterances to help determine context.
   c. language Turing Machine, which maintains and changes states.
   d. language encoder, which turns symbols into neurological impulses and vice versa.

5. Kanwisher claimed that we do not need to be able to speak or understand a language, e.g., English, in order to think because:
   a. a particular region in the brain, rather than the whole brain, is activated for language.
   b. there is no particular region in the brain for language, so language is not internalized uniquely as thought.
   c. babies are able to respond to instructions from their parents even before they can speak.
   d. even people without any functioning language region of the brain are able to perform mathematical tasks.

6. Winston explained that distinguishing external language from inner language matters because:
   a. external language is just a hack to get us understood; inner language is what enables our ability to be symbolic and therefore intelligent.
   b. external language, and the numerous variations of external language, is what distinguishes us from primates; inner language is just an abstraction.
   c. external language and how it changes mimics changes in human thought; inner language is a representation disconnected from how we think.
   d. external language changes too frequently to study; inner language is constant through time.

This page is intentionally blank.

Tear-off sheet

Problem 1 Adaboost

Part A

All Training Points: R, O, M, T, K, S

Weak Classifiers	*Misclassified Training Points*
KP(x)	R, O, M
EP(x)	R, O
CC(x)	M
BP(x)	T, K, S
SS(x)	O, M, T

Tear-off sheet

Problem 2 Bayes

Part A

Part B

P(A)
a

P(B)
b

A	B	P(C|AB)
T	T	c
T	F	d
F	T	e
F	F	f

B	P(D|B)
T	g
F	h

14