6.034 Quiz 4

7 December 2016

Name	
Email	

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

Jake Barnwell	**Michaela Ennis**	**Rebecca Kekelishvili**
Vinny Chakradhar	**Phil Ferguson**	**Nathan Landman**
Alex Charidis	**Stevie Fine**	**Samarth Mohan**
Brian Copeland	**Brittney Johnson**	**Jessica Noss**

Problem	Maximum	Score	Grader
1 - Boosting	50		
2 - Bayes	50		
Total	100		

SRN	6		

There are 10 pages in this quiz, including this one.

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

Problem 1: Football Boosting (50 points)

Part A: Modeling the NFL (33 points)

You want to predict the results of some of this week's key football games. Many factors have an impact on each game, and no individual metric exceeds 65% accuracy in predicting games, so you decide to combine many of these weak classifiers using Adaboost.

In the table below, each training point is a game featuring "Away team @ Home team", and each game is labeled numerically for ease of reference. Each of four weak classifiers predicts a *winner* for each game, either the **Home** team or the **Away** team. The actual winner of each game is listed in the last row of the table.

	Training Points: Seven games from NFL Week 9						
	Game 1 PHI @ NYG	Game 2 DEN @ OAK	Game 3 BUF @ SEA	Game 4 NO @ SF	Game 5 NYJ @ MIA	Game 6 CAR @ LA	Game 7 PIT @ BAL
Fan Poll	Away	Home	Home	Away	Home	Away	Away
Home Team	Home	Home	Home	Home	Home	Home	Home
Better Defense	Away	Away	Home	Away	Home	Home	Home
Better Quarterback	Home	Home	Home	Away	Away	Away	Away
Actual Winner	Home	Home	Home	Away	Home	Away	Home

A1 (3 points) Complete the table below by filling in which training points are misclassified by the fourth weak classifier, *Better Quarterback*:

Weak Classifier	**Misclassified training points**
Fan Poll	1, 7
Home Team	4, 6
Better Defense	1, 2, 6
Better Quarterback	

A2 (30 points) On the next page, perform 2.5 rounds of boosting using these classifiers and training data. In each round, pick the classifier with the **error rate furthest from ½**. Break ties by choosing the classifier that comes earlier in this list: *Fan Poll*, *Home Team*, *Better Defense*, *Better Quarterback*.

In any round, if Adaboost would terminate instead of choosing a classifier, write **NONE** for that round's weak classifier (**h**), then leave all remaining spaces blank.

	Round 1	**Round 2**	**Round 3**
weight 1			
weight 2			
weight 3			
weight 4			
weight 5			
weight 6			
weight 7			
error rate of ***Fan Poll***			
error rate of ***Home Team***			
error rate of ***Better Defense***			
error rate of ***Better Quarterback***			
weak classifier chosen (**h**)			
weak classifier error (**ε**)			
voting power (**α**)			

Show your work for partial credit:

Part B: Predicting the NFL (17 points)

Bruce Headstrong constructs this ensemble classifier:

$$H(x) = 5 \cdot (\textbf{Fan Poll's } vote) + 4 \cdot (\textbf{Home Team's } vote) - 3 \cdot (\textbf{Better Defense's } vote)$$

Classifier	Misclassified
Fan Poll	1, 7
Home Team	4, 6
Better Defense	1, 2, 6

B1 (5 points) Which training points does Bruce's ensemble classifier H(x) misclassify? Circle all that apply, or circle **NONE** if no points are misclassified:

1 2 3 4 5 6 7 NONE

B2 (3 points) Is there some assignment of voting powers (α) that could form a perfect ensemble classifier out of these three weak classifiers?
If so, assign voting powers to the three weak classifiers to form a perfect ensemble classifier. If not, instead circle **CAN'T BE DONE**.

Weak classifier	Voting power
Fan Poll	$\alpha =$
Home Team	$\alpha =$
Better Defense	$\alpha =$

CAN'T BE DONE

B3 (6 points) Now, you want to use Bruce's ensemble classifier to predict three games. Based on the data below, how does Bruce's classifier classify each game? (That is, which team is predicted to win?) Fill in the **Predicted Winner** row by writing **Home** or **Away** in each box. If the answer can't be determined from the available information, instead write **CAN'T TELL**.

	Test Points: Three upcoming games from NFL Week 10		
	Game 8 CIN @ NYG	Game 9 KC @ CAR	Game 10 ATL @ PHI
Fan Poll	Home	Away	Away
Home Team	Home	Home	Home
Better Defense	Home	Away	Home
Better Quarterback	Away	Home	Away
Predicted Winner			

B4 (3 points) The next week, you find out that the *actual winners* of the Week 10 games were:

Game	Game 8	Game 8	Game 10
Actual Winner	Home [NYG]	Away [KC]	Home [PHI]

Given this information, what was the accuracy of Bruce's classifier (i.e. fraction of games correctly predicted) on this test dataset? Fill in the box:

Problem 2: Fantastic Bayes (50 points)

A British wizard has accidentally released some magical creatures at MIT! He offers a reward for safely capturing and returning the creatures. Eager to help, you set off in search of clues!

Part A: Big-Picture Bayes (7 points)

On the chalkboards in Stata, you discover a drawing of an enormous, densely packed Bayes net with variables arranged in a square grid. Upon closer inspection, you find that each variable within the grid has exactly three parents and three children, as shown below, zoomed in:

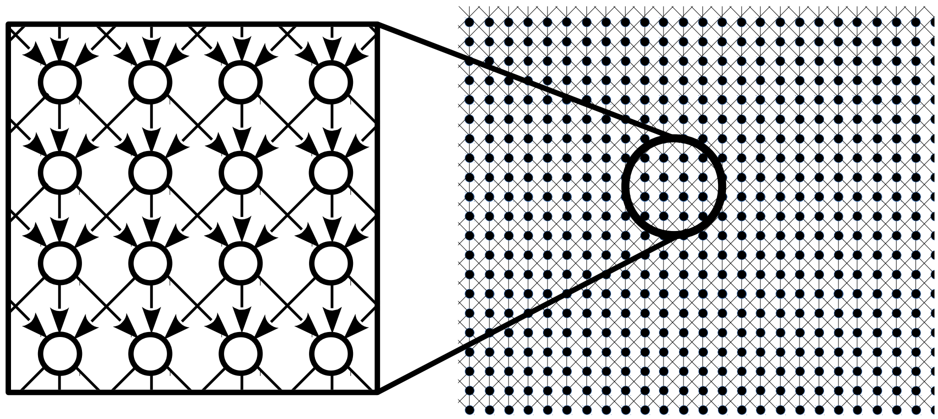

On the next board, you find eight (8) sketches. In each sketch, the shaded region represents some subset of nodes in the enormous Bayes net:

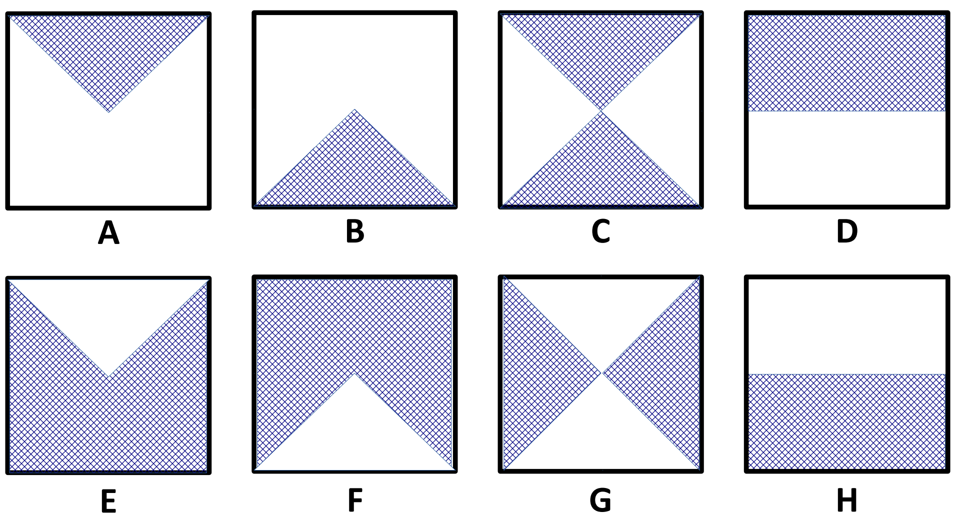

Which one sketch most closely represents…

…the **ancestors** of the center node?

…the **descendants** of the center node?

…the **nodes that are conditionally independent** of the center node, assuming that the center node's parents are given?

Part B: Classifying Fantastic Beasts (25 points)

While walking through Lobby 10, you notice an unfamiliar 4-legged creature running through Killian Court. You know that three 4-legged creatures are still on the loose: a **unicorn**, a **hippogriff**, and a **chimaera**. In your guidebook on Fantastic Beasts, you find the following information on a torn page:

The Bayes net below can be used to distinguish between three 4-legged creatures (classification **Y**) using naïve Bayes classification: a **unicorn**, a **hippogriff**, and a **chimaera**. Some of these creatures **L**ook friendly, and each creature may **C**hase people, in accordance with the probabilities given below:

P(Y=unicorn)	P(Y=hippogriff)	P(Y=chimaera)
0.5	0.4	0.1

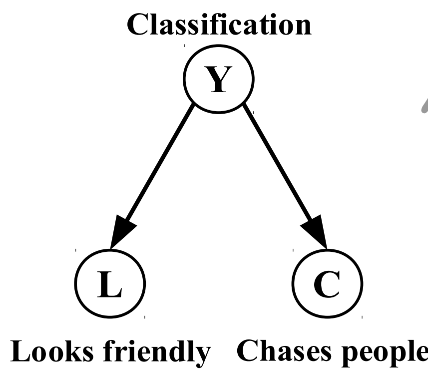

Y	P(L|Y)
unicorn	0.8
hippogriff	0.6
chimaera	0.1

Y	P(C|Y)
unicorn	0.4
hippogriff	0.4
chimaera	0.6

B1 (3 points) Based on the *possible values* of each variable and the *structure* of the Bayes net, what is the **minimum number of parameters** required in the Bayes net?

Show your work for partial credit:

B2 (2 points) You want to classify the unknown creature as a **unicorn**, **hippogriff**, or **chimaera**. Based on the *prior probabilities alone*, circle the **one** most likely classification. If it's impossible to tell, instead circle **CAN'T TELL**:

unicorn **hippogriff** **chimaera** **CAN'T TELL**

B3 (18 points) Rafael Reif comes running into the lobby, looking terrified. You ask him for more information about the creature, such as whether it looked like a horse with a single horn, but he doesn't remember. He replies: *"It was chasing me too fast for me to notice details, but it didn't look friendly!"*

Given that the creature **C**hases people (**C=True**) and does **NOT L**ook friendly (**L=False**), use the Bayes net to determine the likelihood of each classification. **Below, list the classifications (unicorn, hippogriff, chimaera) from *most* to *least* likely**:

Most likely		Least likely

Show your work for partial credit:

B4 (2 points) How would your naïve Bayes classifier classify the creature, given the evidence? (Circle one)

unicorn **hippogriff** **chimaera** **CAN'T TELL**

Part C: Chalkboard Malfunction (18 points)

You enter 10-250 and notice that the **Boards** appear to be sliding up and down by themselves! You realize that this could either be caused by technical **Malfunction**, or by an invisible **Demiguise** (a magical creature) pushing the buttons. Armed with your guidebook on Fantastic Beasts, you draw the following Bayes net:

P(M)
2/10

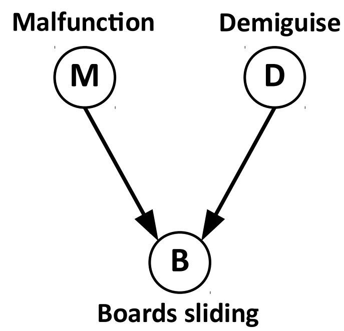

P(D)
1/10

M	D	P(B|MD)
T	T	4/10
T	F	1/10
F	T	3/10
F	F	0

You wonder whether you should be surprised that the boards are sliding without any people present. Without knowing anything about whether there is a **D**emiguise and/or a **M**alfunction, you calculate the probability of the **B**oards would sliding and get **P(B) = 5/100**.

C1 (10 points) Given that the **B**oards are sliding (**B=True**), what is the probability that there is an invisible **D**emiguise near the boards (**D=True**) and that there is **NOT** a **M**alfunction (**M=False**)? Write your answer as a number or numeric expression:

Show your work for partial credit:

C2 (3 points) It's difficult to find an invisible creature, so you wonder whether you should just call MIT FIXIT to report a malfunction. Accordingly, you start computing the probability that there is a **M**alfunction given that the **B**oards are sliding, that is, **P(M|B)**.

But suddenly, a **D**emiguise becomes visible, and you see that it is pushing the buttons! With this new information, the probability of a malfunction is **P(M|BD)**. Given the **D**emiguise's appearance, the probability of a malfunction has: (Circle one)

INCREASED **DECREASED** **STAYED THE SAME**

C3 (2 points) Which effect or principle **best** explains your answer to part C2 above?
Circle the **one** best answer.

1. Conditional Independence
2. Occam's Razor
3. Exhaustion
4. Explaining Away
5. Mutual Exclusion

C4 (3 points) You carefully capture the Demiguise and move away from the buttons. Surprisingly, the boards are *still* sliding! Now, it should be easy to determine the probability of a **M**alfunction, given that the **B**oards are sliding and there is no **D**emiguise pushing the buttons: **P(M=True|B=True, D=False)**. This probability is closest to: (Circle one)

$0 \quad \frac{1}{10} \quad \frac{2}{10} \quad \frac{3}{10} \quad \frac{4}{10} \quad \frac{5}{10} \quad \frac{6}{10} \quad \frac{7}{10} \quad \frac{8}{10} \quad \frac{9}{10} \quad 1$

See reverse for SRN questions →

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. Mansinghka explained how probabilistic programming has been used to:
   1. Improve Wall Street statistical arbitration schemes.
   2. Build cyber defenses against ransomware.
   3. Find characteristics associated with diabetes.
   4. Examine credit card histories to predict likely bankruptcies.
   5. Speed up deep neural net learning.

2. AlphaGo:
   1. Models the way humans play Go.
   2. Picks the move with the highest probability of success using Bayesian inference.
   3. Evaluates all move combinations 20 levels deep using a computing cloud.
   4. Uses Monte Carlo game play in its static evaluator.
   5. Determines each move using k-nearest neighbors in a vast library of human games.

3. Berwick argued that:
   1. All primates have the merge operation.
   2. Merge capability increases with brain size.
   3. Merge is associated with completion of an anatomical loop in the brain.
   4. Recently discovered Neanderthal artifacts demonstrate that they too had the merge capability.
   5. Nim Chimpsky was taught to use merge via training in American Sign Language.

4. Winston described how the Genesis story-understanding system:
   1. Tells stories persuasively by reference to parables found in sacred texts.
   2. Reflects our human tendency to seek explanations.
   3. Deploys neural-net techniques trained on millions of stories to make inferences.
   4. Summarizes stories by retaining only the first *n* and final *m* events.
   5. Uses word co-occurrence statistics to identify concepts such as *revenge*.

5. Pratt emphasized that:
   1. Statistical evidence indicates self-driving cars will cut fatalities per mile to near zero in 10-20 years.
   2. Establishing responsibility for self-driving car accidents will require legal innovation on an unprecedented scale.
   3. The "handoff" problem is about researcher reluctance to transfer development work to production engineers.
   4. Texting and other new distractions are a major reason for increased self-driving car research.
   5. Human drivers are extremely competent when viewed from a fatalities-per-mile perspective.

6. According to Wilson, place fields are:
   1. Neurons in rats' brains that fire based on locations they're thinking about.
   2. Nearby regions in a maze in which rats can localize other rats.
   3. Electrical fields that help rats determine their location.
   4. Heuristic estimates that help rats find their food based on visual stimuli.
   5. Mental maps that rats form to store the shortest path through each maze.