For Quiz #4, of the 73 students who sat the Quiz:
(Due to rounding, numbers shown below are only an approximate estimate.)

Because EEE does not return to you these numbers, in full transparency:

For Total Quiz #4 Score:
"Perfect" (100%): ~22% (16 students)
"A" range (90-100%): ~37% (27 students)
"B" range (80-89%): ~19% (14 students)
"C" range (70-79%): ~14% (10 students)
"D" range (60-69%): ~10% (7 students)
"F" range (<60%): ~21% (15 students)

In the score breakdowns given above, ' "A" range' denotes A+, A, or A-, and so on.

Please, if you are not scoring as highly as you would like to score: attend both lecture and discussion section (and pay attention; do not "multi-task"), attend office hours with the TA or me, schedule an off-hours office meeting with us, review the lecture notes, re-read your book, work the old tests as study guides, and do the homework. Please see the "Study Habits" section on the class website. In short --- OVER-STUDY!!

For each question on Quiz #4, "Zero" gives the percentage of students who received zero, "Partial" gives the percentage who received partial credit, and "Full" gives the percentage who received 100%.

Problem 1
full credit: ~64% (47 students)
partial credit: ~35% (25 students)
zero credit: ~1% (1 student)

Common errors:
- **not building a leaf node when the remaining examples are pure;**
- **not applying information gain correctly.**

Problem 2
full credit: ~23% (17 students)
partial credit: ~74% (54 students)
zero credit: ~3% (2 students)

Common errors:
- **confusing Bayes' rule for the Naive Bayes expression;**
- **for P(Fins=x | Class=y), dividing by the total number of examples instead of dividing by the number of examples that have Class=y;**
- **writing an algebraic expression in 2.c instead of numerical values.**

CS-171, Intro to A.I. — Quiz #4 — Winter Quarter, 2015 — 20 minutes

YOUR NAME AND EMAIL ADDRESS: _____

YOUR ID: _____ ID TO RIGHT: _____ ROW: _____ SEAT: _____

1. (50 pts total) One Fish, Two Fish, Red Fish, Blue Fish. (With apologies to Dr. Suess.) Decision Tree Classifier Learning. You are a robot in the aquarium section of a pet store, and must learn to discriminate Red fish from Blue fish. Unfortunately, your vision sensors are in Black & White, but Red fish have the same gray-scale tone as Blue fish. So, you must learn to discriminate them by body parts. You choose to learn a Decision Tree classifier. You are given the following examples:

This is the same as problem #8 on Final Exam, FQ 2013; except Red fish replaces Oak wood, Blue fish replaces Pine wood, and attribute & value names were changed into a fish theme.

Example	Fins	Tail	Body	Class
Example #1	Thin	Small	Slim	Red
Example #2	Wide	Large	Slim	Red
Example #3	Thin	Large	Slim	Red
Example #4	Wide	Small	Medium	Red
Example #5	Thin	Small	Medium	Blue
Example #6	Wide	Large	Fat	Blue
Example #7	Thin	Large	Fat	Blue
Example #8	Wide	Small	Fat	Blue

See Section 18.3.

If root is Fins:
Thin = RRBB, Wide = RRBB
If root is Tail:
Small = RRBB, Large = RRBB
If root is Body:
Slim=RRR, Medium=RB, Fat=BBB
(R = Red, B = Blue)
Note that you do not need math or a calculator to answer this correctly. Obviously, Fins and Tail do not reduce entropy, while Body does.

1a. (20 pts) Which attribute would information gain choose

Body .

1b. (10 pts) Draw the decision tree that would be constructed by recursively applying information gain to select roots of sub-trees, as in the Decision-Tree-Learning algorithm.

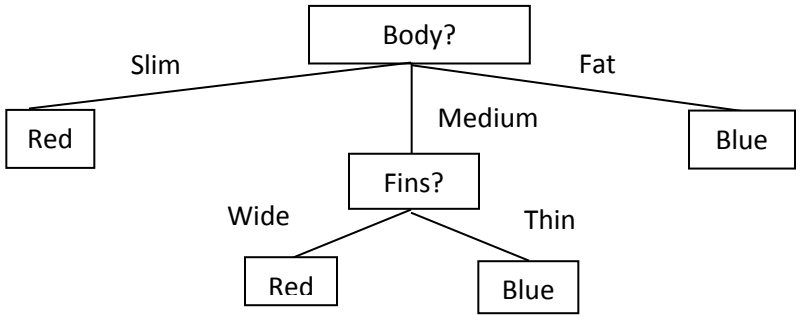

After we choose Body for the root:
*** Tail does not discriminate the two remaining unclassified examples (#4/Red & #5/Blue) because Tail for both is Small.**
*** However, Fins does separate them perfectly, because Fins is Wide for #4/Red and Thin for #5/Blue.**

Classify these new examples as Red or Blue using your decision tree above.

1c. (10 pts) What class is [Fins=Thin, Tail=Small, Body=Fat]? Blue

1d. (10 pts) What class is [Fins=Wide, Tail=Large, Body=Medium]? Red

Full credit if your answers are right for the tree you drew, even if the tree itself is wrong.

****** TURN PAGE OVER AND CONTINUE ON THE OTHER SIDE ******

See Sections 13.5.2 and 20.2.2.

2. (50 pts total) One Fish, Two Fish, Red Fish, Blue Fish. (With apologies to Dr. Suess.) Naïve Bayes Classifier Learning. You are a robot in the aquarium section of a pet store, and must learn to discriminate Red fish from Blue fish. Unfortunately, your vision sensors are in Black & White, but Red fish have the same gray-scale tone as Blue fish. So, you must learn to discriminate them by body parts. You choose to learn a Naïve Bayes classifier. You are given the following examples **(these examples are different from the examples that were given in problem #1, above):**

This is the same as problem #3 on Final Exam, WQ 2012; except Red fish replaces Dog, Blue fish replaces Cat, and attribute & value names were changed into a fish theme.

Example	Fins	Tail	Body	Class
Example #1	Thin	Large	Thin	Red
Example #2	Wide	Small	Thin	Red
Example #3	Wide	Large	Fat	Red
Example #4	Wide	Large	Fat	Red
Example #5	Thin	Small	Thin	Blue
Example #6	Thin	Large	Fat	Blue
Example #7	Wide	Small	Fat	Blue
Example #8	Thin	Small	Thin	Blue

Unfortunately, your textbook uses an inconsistent notation to refer to values of attributes. In Chap. 13, values of attributes (= random variables) are lower-case (see Section 13.2.2). In Chap. 18, values of attributes are upper-case (see Fig. 18.3). Here, since we are in the machine learning part of the course, we will follow Chap. 18 and use upper-case values of attributes. Please do not be confused.

…es' rule allows you to rewrite the conditional probability of the class given the attributes as the conditional probability of the attributes given the class. As usual, α is … that makes the likelihoods (unnormalized probabilities) sum to one. Thus, we … repeated denominator **P**(Fins, Tail, Body), because it is constant for all class… we rewrite: **P**(Class | Fins, Tail, Body) = α **P**(Fins, Tail, Body | Class) P(Class…

The probabilities in problem 2b are obtained by counting examples in the training set. E.g., P(Class=Red)=4/8=1/2 because 4 of the 8 examples have Class=Red. E.g., P(Fins=Thin|Class=Red)=1/4 because 1 of the 4 examples with Class=Red also has Fins=Thin.

2a. (10 pts) Now assume that the attributes (Fins, Tail, and Body) are condi… given the Class. Rewrite the expression above, using this assumption of con… (i.e., rewrite it as a Naïve Bayes Classifier expression).

α **P(Fins, Tail, Body | Class) P(Class) =** α **P**(Fins | Class) P(Tail | Class) P(Body | Class) P(Class)

2b. (20 pts total; −2 for each wrong answer, but not negative) Fill in num… following expressions. Leave your answers as simplified common fractions

P(Class=Red)= 1/2 P(Class=Blue)= 1/2

P(Fins=Thin | Class=Red)= 1/4 P(Fins=Thin | Class=Blue)= 3/4

P(Fins=Wide | Class=Red)= 3/4 P(Fins=Wide | Class=Blue)= 1/4

P(Tail=Large | Class=Red)= 3/4 P(Tail=Large | Class=Blue)= 1/4

P(Tail=Small | Class=Red)= 1/4 P(Tail=Small | Class=Blue)= 3/4

P(Body=Thin | Class=Red)= 1/2 P(Body=Thin | Class=Blue)= 1/2

P(Body=Fat | Class=Red)= 1/2 P(Body=Fat | Class=Blue)= 1/2

2c. (20 pts total, 10 pts each) Consider a new example **(Fins=Wide ^ Tail=Large ^ Body=Thin)**. Write these class probabilities as the product of α and common fractions from above. You do not need to produce an actual final number; only an expression that will evaluate to the right answer.

2.c.i (10 pts) P(Class=Red | Fins=Wide ^ Tail=Large ^ Body=Thin)

= α(3/4)(3/4)(1/2)(1/2) **(=9/10)**

2.c.ii (10 pts) P(Class=Blue | Fins=Wide ^ Tail=Large ^ Body=Thin)

= α(1/4)(1/4)(1/2)(1/2) **(=1/10)**

You are not obliged to provide the (red) "(=9/10)" and "(=1/10)" evaluations; only the fractional products in black that precede them in the answer. The final normalized probabilities (in red) are only for your information in seeing it work.