HW2 – COGS001 Fall 2006

Show all the steps in your calculations.
Avoid rounding intermediate results, to reduce cumulative numerical errors.
Notation: ^ denotes exponentiation.

1. a. Connectionist modeling, as described by Dawson, is sometimes used to develop theories about cognition at the procedural/algorithmic level.
i) Identify one advantage of using connectionist models for this purpose. [One sentence.]
ii) Identify one disadvantage compared to classical architectures (such as Turing machines and other symbol systems). [One sentence.]

b. i) Identify one reason for believing connectionist models provide poor theories at the implementation level. [One sentence.]
ii) Identify a counterargument to b(i), in support of connectionism. [One sentence.]

2. I wake up to the sound of my car alarm going off. There are various possible events that could trigger my car alarm, namely disturbance by a burglar, by a passing car, or by a thunderstorm. Or it could be something else entirely. The base rates for these events occurring are detailed below. Assume that these events are mutually exclusive, so no two can occur simultaneously.

B - burglar	p(B) = 0.02
C - passing car	p(C) = 0.5
T - thunderstorm	p(T) = 0.18
N - none of above	p(N) = ?

The alarm manufacturer states the following likelihoods for my car alarm being triggered given each event.

p(A|B) = 0.95
p(A|C) = 0.0005
p(A|T) = 0.01
p(A|N) = 0.001

(i) Calculate the p(N). This is the prior probability that none of B, C or T occurs.

(ii) Calculate:
a. the probability that my car is burgled *and* the alarm sounds, i.e. p(A,B);
b. the probability that another car passes *and* the alarm sounds, i.e. p(A,C);
c. the probability that storm occurs *and* the alarm sounds, i.e. p(A,T);
d. the probability that none of B/C/T occur *and* the alarm sounds, i.e. p(A,N).

(iii) Using your answers to (ii), calculate the probability of my car alarm being triggered in general, i.e. p(A).

(iv) Given that my car alarm is currently going off, what is the probability that:
a. there is a burglar?
b. there is a passing car?
c. there is a thunderstorm?

3. Fred the Frequentist is tested for a disease. This type of test is known to be accurate 97% of the time (i.e., 97% of people who carry the disease will receive a positive result and 97% of people who do not carry the disease receive a negative result). Demographics indicate that 1 in 500 people of Fred's age, gender and background have this disease.

We denote Fred's disease-status and his test result by the variables h and t respectively.
Fred may either have (h=1) or not have (h=0) the disease.
His test result may be either positive (t=1) or negative (t=0).

a. State the values of the following six probabilities: p(h=0), p(h=1), p(t=1|h=1), p(t=1|h=0), p(t=0|h=1), p(t=0|h=0).
b. Write p(t=1) in terms of p(h=0), p(h=1), p(t=1|h=1), p(t=1|h=0) and calculate its value.
c. Before taking the test Fred feels happy. At this point, what are the chances he has the disease?
d. Fred tests positive for the disease. Suddenly, he feels he is doomed, since he believes his chances of having the disease are now 97%. But, having studied Bayes' Rule in COGS001, you know better! Given that Fred has tested positive, what is the probability that Fred has the disease?

4. Google is scanning millions of books from the world's libraries into machine-readable files. I would like to automatically classify electronic books into one of three genres: Biography (B), Mystery (M), or Historical (H).

My sample includes only these three genres, in the following proportions:

P(B) = .26;	P(M) = .32;	P(H) = .42.

The likelihood of seeing a particular word in given book depends on the book's genre:

P("shadow"|B) = .3;	P("shadow"|M) = .6;	P("shadow"|H) = .15;
P("Marilyn"|B) = .9;	P("Marilyn"|M) = .2;	P("Marilyn"|H) = .3;
P("war"|B) = .3;	P("war"|M) = .1;	P("war"|H) = .8.

Various common words ('the,' 'in,' 'was,' 'her,' 'he,' 'to,' 'had') are equally likely in all genres.

a. Consider classifying a book using just one word. I open a book and see "shadow".
(i) Do I need to know the probability P("shadow") in order to compare the posterior probabilities of the 3 genres? Explain in one sentence.
(ii) Which genre maximizes the posterior probability P(genre|"shadow")? Show your calculations.

b. Consider classifying a book using snippets of text. Which genre would the Naïve Bayes classifier assign to the following three snippets of text? Show your calculations.

(i) "Marilyn was in her shadow"
(ii) "he had to shadow Marilyn"
(iii) "the shadow of the war"

c. Does Naïve Bayes explicitly model the different meanings of the word "shadow" in snippets (i)-(iii)? Explain briefly what your answer implies about this classifier's fallibility.

d. I expand my sample. Now there are N genres and the vocabulary is a set of W possible words.
I consider three different classifier models for this task:
- ***Naïve Bayes model***: as used above.
- ***Bag of Words model***: a likelihood is assigned to the appearance of every possible subset of the W words in the vocabulary, given each of the N genres. Note: there are 2^W possible subsets of a set with W elements.
- ***Page of Words model***: a likelihood is assigned to the appearance of every possible permutation (ordering) of words on a page, given each of the N genres. Note: assume there are D words per page, so there are W^D possible permutations of words on a page. (A permutation is an ordering. For example if you had 2 words (A,B) in your vocabulary and each page is 3 words long, then the 2^3 possible permutations are: AAA, AAB, ABB, ABA, BBB, BBA, BAB, BAA.)

(i) In terms of N and W, write an expression for the number of parameters required for the Naïve Bayes classifier. Is your answer linear or exponential in N and W?

(ii) In terms of N and W, write and expression for the number of parameters required for the Bag of Words classifier. Is your answer linear or exponential in N and W?

(iii) In terms of N, W and D, write and expression for the number of parameters required for the Page of Words classifier. Is your answer linear or exponential in N and D?

(iv) In 2 words, what does Naïve Bayes assume about the words, to overcome the curse of dimensionality?

5. You find yourself on an island, unable to speak the native language.

Animals on this island can be described according to three concepts:

Species - out of 50 possible species, with *each species equally represented*;
Color - out of 10 possible colors, with *each color equally represented*;
Gender - out of 2 possible genders, with *each gender equally represented*.

In fact, the island is keen on egalitarianism, such that *every "species, color, gender" combination is equally represented*. This means that both genders of every species exist equally in all ten colors! Moreover, this means that species, color and gender are *conditionally independent*!

An inhabitant points to a passing animal and says a word. Your goal is to model what the word may mean, using Bayesian inference (similar process to the "numbers game" in lectures).

There are 3 classes of hypotheses for a word, and any word must satisfy one and only one class:

H=Hs : word refers to the species
H=Hc : word refers to the color
H=Hg : word refers to the gender

You model the three classes to be equally probable – that is to say, you believe the inhabitant is equally likely to describe either the species, or the color, or the gender. Whichever concept he chooses to describe, he will describe it accurately.

First we will look at the prior probabilities in this scenario, before you witness any data.

a. Intuitively, under which hypothesis class would data (a word) be "more of a coincidence", in the sense described in lectures?

b. What are the prior probabilities of each hypothesis class?
i.e., what is: (i) P(H=Hs); (ii) P(H=Hc); (iii) P(H=Hg)?

c. Recall that every "species, color, gender" combination is equally represented. Given each hypothesis class, what is the probability that a particular category within that hypothesis class is mentioned? So, for example, what is:
(i) P(s=kangaroo | H=Hs) - i.e., the probability that the species described is a kangaroo, *given that* the word describes species;
(ii) P(c=pink | H=Hc) - i.e., the probability that the color described is pink, *given that* the word describes color;
(iii) P(g=male | H=Hg) - i.e., the probability that the gender described is male, *given that* word describes gender?

d. Using your answers to (b) and (c), what are the prior probabilities of each hypothesis?
So, for example, what is:
(i) P(s=elephant, H=Hs) - i.e., the probability that the word describes species *and* the species described is an elephant;
(ii) P(c=purple, H=Hc) - i.e., the probability that the word describes color *and* the color described is purple;
(iii) P(g=female, H=Hg) - i.e., the probability that the word describes gender *and* the gender described is female?

e. Comparing just prior probabilities of each hypothesis, is a word more likely to mean rabbit or male or white?

Question 5 continued…

Now we will consider the implications of the data you observe. Initially your data consists of one observation: the inhabitant pointed to a male white rabbit (MWR) and declared, "kikibouba!"

f. Recall that the species, color and gender of animals on this island are conditionally independent. What is the likelihood that the animal referred to is a MWR, *given that* "kikibouba" means:
(i) male – i.e., what is P(MWR | g=male, H=Hg);
(ii) white – i.e., what is P(MWR | c=white, H=Hc);
(iii) rabbit – i.e., what is P(MWR | s=rabbit, H=Hs)?
Which of these three hypotheses maximizes the likelihood of the data?

g. Using marginalization over all possible hypotheses, calculate P(MWR) – i.e., the probability that a MWR is referred to in general on the island.

h. Given that the animal referred to is a MWR, what is the posterior probability that "kikibouba" means:
(i) male – i.e., what is P(g=male, H=Hg | MWR);
(ii) white – i.e., what is P(c=white, H=Hc | MWR);
(iii) rabbit – i.e., what is P(s=rabbit, H=Hs | MWR);
(iv) female – i.e., what is P(g=female, H=Hg | MWR);
(v) pink – i.e., what is P(c=pink, H=Hc | MWR);
(vi) hippo – i.e., what is P(s=hippo, H=Hs | MWR)?
What can you conclude after one observation?

You make a second observation: a male brown rabbit (MBR) is also referred to as "kikibouba".

Note that, after your first observation, you know more about "kikibouba" than you did at the outset. So, when dealing with new observations regarding "kikibouba", ***you should use the posterior probabilities that you calculated in (h) as your new priors***.

For instance, in (h) you calculated P(s=rabbit, H=Hs | MWR). This will become the value for P(s=rabbit, H=Hs) from now on. The posterior probability after the first observation is the prior for the second observation!

i. Using marginalization (*for which you will be using the posterior probabilities from part (h) as your priors)* over all possible hypotheses, calculate P(MBR).

j. Given that "kikibouba" also refers to a MBR, what is the posterior probability that "kikibouba" means
(i) male – i.e., what is P(g=male, H=Hg | MBR);
(ii) white – i.e., what is P(c=white, H=Hc | MBR);
(iii) rabbit – i.e., what is P(s=rabbit, H=Hs | MBR)?
Again, *use the posterior probabilities from part (h) as your priors here*, to account for already having observed that "kikibouba" also refers to a MWR.

6. Subjects in an experiment read transcripts of brief interviews. From the transcripts they learn general facts about the interviewees' personalities. After reading each one, they are asked to judge whether the interviewee would prefer to eat out at Lemongrass or Qdoba.

In addition, subjects are told the following: *"75% of these interviews were conducted at Lemongrass and 25% were conducted at Qdoba. Every interviewee favored the restaurant in which they were eating, so 75% of the sample likes Lemongrass and 25% likes Qdoba."*

a. In the terminology of Tversky and Kahneman, is the base rate (.75 vs. .25) causal or incidental?

b. What proportion of the interviewees should subjects assign to each restaurant, if they are taking base rates into account?

c. How could you rewrite the text in quotation marks above to change the *type* of base rate?
(That is, if it's an incidental base rate, rewrite to make it a causal base rate. If it's a causal base rate, rewrite to make it an incidental base rate. Do not change the value of the base rate.)

d. Would your modification increase or decrease the degree to which subjects take the base rate into account? (That is, would it bring them closer to, or further from, the proportion in part (b)?)

e. According to Tversky and Kahneman's rationale, why would it have this effect?