CMSC 478 Machine Learning - Spring 2019
Homework Assignment 4
Due at the start of class on March 14^{th}

1. (Double counting the evidence) Consider a problem in which the class label $y \in \{T, F\}$ and each training example X has 2 binary attributes $X_1, X_2 \in \{T, F\}$.

   Let the class prior be $p(Y = T) = 0.5$ and $p(X_1 = T|Y = T) = 0.8$ and $p(X_2 = T|Y = T) = 0.5$. Likewise, $p(X_1 = F|Y = F) = 0.7$ and $p(X_2 = F|Y = F) = 0.9$. Attribute X_1 provides slightly stronger evidence about the class label than X_2.

   Note: From this information we can infer the following:

   - $p(Y = F) = 0.5$ - because $p(Y = T) + p(Y = F) = 1$
   - $p(X_1 = F|Y = T) = 0.2$ - because $p(X_1 = T|Y = T) + p(X_1 = F|Y = T) = 1$
   - $p(X_2 = F|Y = T) = 0.5$
   - $p(X_1 = T|Y = F) = 0.3$
   - $p(X_2 = T|Y = F) = 0.1$

   - Assume X_1 and X_2 are truly independent given Y. Write down the naive Bayes decision rule. That is, write down the formulas for the two quantities that you compare to decide if the class label is T or F.
   - What is the expected error rate of naive Bayes if it uses only attribute X_1? What if it uses only X_2?

     The expected error rate is the probability that each class generates an observation where the decision rule is incorrect. If Y is the true class label, let $\hat{Y}(X_1, X_2)$ be the predicted class label. Then the expected error rate is $p(X_1, X_2, Y|Y \neq \hat{Y}(X_1, X_2))$.

     Here is how you compute the expected error rate when you just use X_1. Note that X_1 can be either T or F. The table below shows for each possible observed value of X_1, the quantities that must be computed and compared to determine the predicted class label. When $X_1 = T$, the prediction is $Y = T$. When $X_1 = F$, the prediction is $Y = F$.

     The expected error rate is the summed probabilities of those cases in which the predictions are wrong, which is $p(X_1 = T, Y = F) + p(X_1 = F, Y = T) = 0.15 + 0.1 = 0.25$.

X_1	$p(X_1\|Y=T)p(Y=T)$	$p(X_1\|Y=F)p(Y=F)$	prediction
T	$0.8 * 0.5 = 0.4$	$0.3 * 0.5 = 0.15$	$Y = T$
F	$0.2 * 0.5 = 0.1$	$0.7 * 0.5 = 0.35$	$Y = F$

- Show that if naive Bayes uses both attributes, X_1 and X_2, the error rate is 0.235, which is better than if using only a single attribute (X_1 or X_2). If this case the table will have 4 rows.
- Now suppose that we create new attribute X_3 which is an exact copy of X_2. So for every training example, attributes X_2 and X_3 have the same value. What is the expected error of naive Bayes now? This table will have 4 rows rather than 8 because cases in which $X_2 \neq X_3$ cannot happen.
- Explain what is happening with naive Bayes? Does logistic regression suffer from the same problem? Explain why.