9. (20 points) In a particular domain, we are able to observe three real-valued input variables X_1, X_2, and X_3 and want to predict a single binary output variable Y (which can have values 0 or 1). We know the functional forms for the input variables are all uniform distributions, namely: $X_1 \sim \text{Uni}(a_1, b_1)$, $X_2 \sim \text{Uni}(a_2, b_2)$, and $X_3 \sim \text{Uni}(a_3, b_3)$, but we are not given the values of the parameters a_1, b_1, a_2, b_2, a_3 or b_3. We are, however, given the following data set of 8 training instances:

X_1	X_2	X_3	Y
0.1	0.8	0.4	0
0.7	0.6	0.1	0
0.3	0.7	0.2	0
0.4	0.4	0.6	0
0.8	0.2	0.5	1
0.5	0.7	0.8	1
0.9	0.4	0.7	1
0.6	0.6	0.4	1

a. (10 points) Use Maximum Likelihood Estimators to estimate the parameters a_1, b_1, a_2, b_2, a_3 and b_3 in the case where Y = 0 as well as the case Y = 1. (I.e., estimate the distribution $P(X_i \mid Y)$ for i = 1, 2, and 3). Note that the parameter values for a_1, b_1, a_2, b_2, a_3 and b_3 may be different when Y = 0 versus when Y = 1.

b. (10 points) You are given the following 3 testing instances, numbered 1, 2 and 3. (Note that the testing instances do not have output variable Y specified).

	X_1	X_2	X_3
test instance 1	0.5	0.6	0.4
test instance 2	0.7	0.7	0.7
test instance 3	0.5	0.4	0.3

Using the Naive Bayes assumption and your probability estimates from part (a), predict the output variable Y for *each* instance (you should have 3 predictions). Show how you derived the prediction by showing the computations you made for each test instance.