Homework 4

Below are some movie reviews in two classes: positive (+) and negative (-).

Number of unique words in the Vocabulary, |V| = 22

	Category C	Documents D
Training Example	-	just plain boring film
Training Example	-	entirely predictable and lacks energy
Training Example	-	no surprises and very few laughs
Training Example	-	far from fantastic
Training Example	+	very powerful drama
Training Example	+	the most fun film
Test Example 1	??	most powerful drama
Test Example 2	??	a fun fantastic film

We represent a document d as a bag-of-words and apply the naive Bayes classifier to estimate the correct class $\hat{c}$ of d:

$$\hat{c} = \arg\max_{c \in C} P(c|d) = \arg\max_{c \in C} \frac{P(d|c)P(c)}{P(d)} = \arg\max_{c \in C} P(d|c)P(c) = \arg\max_{c \in C} P(c) \prod_i P(w_i|c)$$

where i is the index to go through every word in the document d.

We have learned in class how to estimate the probabilities $P(c)$ and $P(w_i|c)$:

$$\hat{P}(c) = \frac{N_c}{N_{doc}} \quad \text{and} \quad \hat{P}(w_i|c) = \frac{count(w_i, c)}{\sum_{w \in V} count(w, c)}$$

1. Using the above equations, predict a label for both Test Examples. (Predict "-" in case of a tie.)

2. Use "Add-One" smoothing below, and recalculate the predictions. Which is better?

$$\hat{P}(w_i|c) = \frac{count(w_i, c) + 1}{\sum_{w \in V} (count(w, c) + 1)} = \frac{count(w_i, c) + 1}{(\sum_{w \in V} count(w, c)) + |V|}$$