Text Classification: Exam practice

> ⓘ **Note**
>
> This document gives *examples* for tasks similar to those that will appear on the digital written exam. The solutions are provided at the end of this document, in case you first want to try solving the tasks yourself. This is not meant as an exhaustive list; the exam may also feature other types of tasks/questions than the ones shown here.

Task 1.1

Here is a collection of documents and their gold-standard class labels:

document	class
awesome	positive
awesome great	positive
great bad	neutral
bad	negative

Use Maximum Likelihood Estimation (MLE) with add-one smoothing to estimate the class probabilities and word probabilities of a Naive Bayes classifier from the given document collection. Assume that the vocabulary consists of the set of all words occurring in the documents.

	class probability	word probability for *awesome*	word probability for *great*	word probability for *bad*
positive				
neutral				
negative				

Task 1.2

A Naive Bayes classifier has been trained on classifying whether an e-mail is spam or not, resulting in the following class probabilities and (some of the) word probabilities:

	class probability	word probability for *money*	word probability for *urgent*	word probability for *tomorrow*
spam	$\frac{1}{3}$	$\frac{45}{500}$	$\frac{55}{500}$	$\frac{1}{500}$
not spam	$\frac{2}{3}$	$\frac{15}{800}$	$\frac{30}{800}$	$\frac{75}{800}$

Based on the probability values above, compute the class-specific scores that the Naive Bayes classifier uses to predict the class for the following document:

tomorrow urgent

	estimated probability
spam	
not spam	

Task 1.3

Given the following counts of sentiment class labels in a document collection, what is the accuracy of the most frequent class baseline on the test data?

	positive	neutral	negative
training data	250	105	225
test data	85	55	90

Task 1.4

Below is a confusion matrix from evaluating a classifier that predicts which category of a newspaper an article was published under. Rows correspond to gold-standard class labels, while columns correspond to predicted class labels; for example, the **highlighted** cell contains the number of times the classifier predicted *"lifestyle"* where the gold-standard class was *"culture"*.

	news	sports	culture	lifestyle
news	505	0	2	4
sports	0	428	15	8
culture	1	4	320	45
lifestyle	10	15	72	155

Based on the confusion matrix above, compute the following evaluation metrics:

precision with respect to *"sports"*	recall with respect to *"lifestyle"*

> 🗨 **Solutions begin on the next page!**

Solutions

Task 1.1

	class probability	word probability for *awesome*	word probability for *great*	word probability for *bad*
positive	$\frac{2}{4}$	$\frac{2+1}{3+1\times3}$	$\frac{1+1}{3+1\times3}$	$\frac{0+1}{3+1\times3}$
neutral	$\frac{1}{4}$	$\frac{0+1}{2+1\times3}$	$\frac{1+1}{2+1\times3}$	$\frac{1+1}{2+1\times3}$
negative	$\frac{1}{4}$	$\frac{0+1}{1+1\times3}$	$\frac{0+1}{1+1\times3}$	$\frac{1+1}{1+1\times3}$

> ⓘ **Note**
>
> We give the unsimplified fractions here to make it clearer where the numbers are coming from: for example, the class probability for the positive class is $\frac{2}{4}$ because there are *two* documents of positive class out of *four* documents in total. In the exam, it doesn't matter if you enter $\frac{2}{4}$, $\frac{1}{2}$, or 0.5 as your solution, as long as it's the same number.

Task 1.2

	estimated probability
spam	$\frac{1}{3} \times \frac{1}{500} \times \frac{55}{500}$
not spam	$\frac{2}{3} \times \frac{75}{800} \times \frac{30}{800}$

> ⓘ **Note**
>
> In the exam, it is fine to enter the equations as shown above; you don't need to simplify them or compute the final result.

Task 1.3

$$\frac{85}{85 + 55 + 90}$$

Task 1.4

precision with respect to *"sports"*	recall with respect to *"lifestyle"*
$\frac{428}{0+428+4+15}$	$\frac{155}{10+15+72+155}$