

School of Computing and Information Systems
The University of Melbourne
COMP90042 WEB SEARCH AND TEXT ANALYSIS (Semester 1, 2019)

Workshop exercises: Week 5

Discussion

1. What is **text classification**? Give some examples.
 - (a) Why is text classification generally a difficult problem? What are some hurdles that need to be overcome?
 - (b) Consider some (supervised) text classification problem, and discuss whether the following (supervised) machine learning models would be suitable:
 - i. k -Nearest Neighbour using Euclidean distance
 - ii. k -Nearest Neighbour using Cosine similarity
 - iii. Decision Trees using Information Gain
 - iv. Naive Bayes
 - v. Logistic Regression
 - vi. Support Vector Machines
2. For the following “corpus” of two documents:
 1. how much wood would a wood chuck chuck if a wood chuck would chuck wood
 2. a wood chuck would chuck the wood he could chuck if a wood chuck would chuck wood
 - (a) Which of the following sentences: a wood could chuck; wood would a chuck; is more probable, according to:
 - i. An unsmoothed uni-gram language model?
 - ii. A uni-gram language model, with Laplacian (“add-one”) smoothing?
 - iii. An unsmoothed bi-gram language model?
 - iv. A bi-gram language model, with Laplacian smoothing?
 - v. An unsmoothed tri-gram language model?
 - vi. A tri-gram language model, with Laplacian smoothing?
3. What does **back-off** mean, in the context of smoothing a language model? What does **interpolation** refer to?

Programming

1. In the `WSTA_N7_text_classification` notebook, observe how different tokenisation regimes alter the text classification performance of the various classifiers on the given Reuters dataset problem.
 - (a) Alter the tokenisation strategy so that it incorporates other stages, for example, punctuation, or stemming/lemmatisation.
 - (b) Does performance increase or decrease? Are some classifiers affected more than others? Why do you think that is?

2. Using the iPython notebook `WSTA_N8_n-gram_language_models`, randomly generate some sentences based on the bi-gram models of the Gutenberg corpus and the Penn Treebank. What do you notice about these sentences? Are there any sentences which might get returned for both corpora? Why?
3. Find a sentence with a higher probability than *revenue increased last quarter.*, according to:
 - (a) The Gutenberg corpus, using bi-grams smoothed with Laplacian smoothing
 - (b) The Gutenberg corpus, using bi-grams smoothed with Interpolation
 - (c) The Penn Treebank corpus, using bi-grams and Laplacian smoothing
 - (d) The Penn Treebank corpus, using bi-grams and Interpolation
4. Find the perplexity of the above (smoothed) language models for a number of sentences. Why does Interpolation generally have better perplexity?

Catch-up

- What is a **language model**? What is an **n -gram language model**? Why are language models important?
- What do **uni-gram**, **bi-gram**, **tri-gram**, etc. signify?
- Why is **smoothing** important?
- Why do we usually use **log probabilities** when finding the probability of a sentence according to an n -gram language model?
- How might one evaluate a language model?

Get ahead

- Adjust the `WSTA_N7_text_classification` iPython notebook, so that the supervised machine learning model attempts to solve the **multi-class** problem, rather than the **single-class** problem (for `acq`). Does your assessment of the relative utility of the given classifiers change?
- Using the (short) “corpus” from Discussion Q1, generate all of the sentences of length 3. Choose an n -gram language model, and find the most probable sentence. What about length 4? 5? 6? What do you notice about these sentences? Does smoothing (or not) change this?
- Modify the iPython notebook so that it uses back-off smoothing. How does this change the probability of the given sentence? Why? Is the perplexity of this model better than Laplacian smoothing? Interpolation? Why?
- Perform the Programming experiments above using different corpora.