

Natural Language Processing

Assignment 4: Semantics

1 Introduction

This assignment involves material from lectures 11 and 12. You should have watched the relevant videos, read the relevant chapters in the textbook and made a serious attempt at completing the relevant labs before you attempt this assignment. If you feel you have done that and still find the instructions unclear, you are welcome to email the course teachers and/or go to office hours to ask for help. The assignment is split into 2 sections, one about lexical semantics, one about semantic role labelling. Each section is worth equally much. We expect between half a page and a page for each section. Please do not submit more than 4 pages overall. Your answers for each section should be self-contained.

2 Lexical Semantics: Error analysis

In this exercise you will perform an error analysis which is based on the results you got from training the classifiers in Lab 11. In Lab 11, you have already analyzed the performance of a classifier. However, in addition to that it is useful to perform an error analysis. By providing additional arguments to `wsd_classifier()`, you can get a confusion matrix as well as a printout of all the errors.

```
>>> nb_hard = wc.wsd_classifier(nltk.NaiveBayesClassifier.train, 'hard.pos', wc.wsd_word_features, distance=3, confusion_matrix=True, log=True)
Reading data...
Senses: HARD1 HARD2 HARD3
Training classifier...
Testing classifier...
Accuracy: 0.8016
Writing errors to errors.txt
  | H H H |
  | A A A |
  | R R R |
  | D D D |
  | 1 2 3 |
-----+-----+
HARD1 |<692> 7 3 |
HARD2 | 84 <2> 2 |
HARD3 | 75 1 <1> |
-----+-----+
(row = reference; col = test)
```

The relevant arguments are `confusion_matrix=True` and `log=True`, but note that you must also insert the argument `distance=3` to get the right interpretation. The errors are printed to an external file called `errors.txt`. Do an error analysis of the **best performing classifier for each target word**, answering the following questions:

1. Using the confusion matrix, identify which sense is the hardest one for the model to estimate. Look in `errors.txt` for examples where that hardest word sense is the correct label. Do you see any patterns or systematic errors? If so, can you think of a way to adapt the feature representation to improve the model? [ca. 1/2 page of discussion]
2. Get familiar with the concept of Precision, Recall and F-Measure (coursebook, Chapter 13.5.3) and calculate them based on the contingency table of the best performing classifier for one of the words. Then, discuss the (possible) advantages of Precision, Recall and F-measure as opposed to Accuracy when you want to improve the performance of a tool (in our case: a classifier) [ca. 1/2 page]

3 Semantic Role Labeling

In Lab 12, you annotated some data and discussed with your classmates. Describe this experience and discuss inter-annotator agreement. Give examples of cases where agreement was good and where it was bad. Look up measures of inter-annotator agreement on the internet and calculate one of them.

[at least 1 page for this section]

4 Grading Criteria

To pass the assignment, you must meet all the basic criteria on all subparts of the assignment. To get VG, you must in addition meet some of the additional criteria for most of subparts.

Basic Criteria

- Answers are given in understandable English.
- Answers are stated clearly and coherently.
- Answers are essentially correct.

Additional Criteria

- Answers are well motivated.
- Answers are well illustrated.
- Answers reveal extensive knowledge of the textbook chapter(s).

5 Submit the Assignment

Submit your assignment as a pdf file named `firstname.lastname.assignment_4.pdf`. It should follow the style and margins given in the example submission even if not created with LaTeX. The submission is due on *studentportalen* before Friday December 14th at 20h00. Later submissions will be considered failed submissions and assessed after the final re-submission deadline on January 11th.