

# Distinguishing Opinion from News

Katherine Busch

## *Abstract*

Newspapers have separate sections for opinion articles and news articles. The goal of this project is to classify articles as opinion versus news and also to do analysis of the results to figure out the factors that distinguish the two. Preliminary results show that in this is possible with unigram features in an SVM with F1 of .90.

## *Introduction*

This project focuses on subjectivity classification for news articles. Much prior work on subjectivity has focused on distinguishing positive and negative sentiment (for instance, in product reviews) or classifying phrases or clauses as subjective (Liu, 2008). Here we attempt to distinguish entire articles as reporting news or expressing opinion. The task is related but has some key differences. For instance, review-type sentiment analysis often relies on pre-made lexicons or focuses on classifying words as positive or negative (Toprak and Gurevych, 2009; Turney and Littman, 2010; Potts). Words associated with positivity and negativity are not necessarily those associated with editorials and opinion pieces in which authors pose sophisticated arguments about current events, policies, etc. One goal of the project was to gain a lexical understanding of words that can distinguish the two categories, and thus be able to generate a lexicon similar those already existing for sentiment analysis of reviews that would work for articles.

## *Prior Work*

There has been thorough research into document classification. "Machine Learning in Automated Text Categorization" (Sebastiani, 2003) provides an overview of work up to 2002. Within the area of subjectivity/sentiment analysis there is also a wide variety of work. Pang and Lee give an overview of the field of subjectivity (Pang and Lee, 2008). Liu defines many different problems within the field including *sentiment and subjectivity classification*:

- (1) classifying an opinionated document as expressing a positive or negative opinion, and (2) classifying a sentence or a clause of the sentence as subjective or objective

(Liu, 2010)

Liu also gives an overview of the field thus far from a teaching perspective. Turney and Littman provide a method for sentiment for particular words based on their context (Turney and Littman, 2003). Yu and Hatzivassiloglou specifically address distinguishing opinion from news using a Naive Bayes classifier and are able to achieve very high results (Yu and Hatzivassiloglou, 2003).

## Data

I use two datasets, both consisting of articles from the *New York Times*. The primary dataset consists of 15 over the course of the 7 years up to and including 2012. For comparison, I also test on a dataset of articles from October and November 2012 in which news events are covered repeatedly. The data was collected by scraping the *New York Times* website. The first set includes 10 news articles and 5 opinion articles/year arbitrarily selected. The second includes the entire world and United States news sections and entire opinion sections for several days in the past months.

## Results

### October-November 2012

Learning Algorithm	F1
Multinomial Naive Bayes (no smoothing)	.46
Multinomial Naive Bayes (Laplace smoothing)	.87
SVM: Unigram counts	.83

Table 1: F1 for small time period dataset

### 2006-2012

Learning Algorithm	F1
SVM: unigram counts	.85
SVM: unigram counts with stemming	.88
SVM: TFIDF	.70
SVM with PoS tags counts (32 features)	.67
SVM with PoS and unigram	.85
<b>SVM with top 600 features</b>	<b>.90</b>

Table 2: F1 for large time period dataset

## Analysis

I focus on the results for the mixed years dataset and only use the small time period dataset for comparison in the *Language results* section. Overall, our classifier achieved high precision and recall for the test set with the best F1 score of .9. Of the four misclassified articles using stemming and all features for the mixed years, two were reviews (one in travel and one in technology) that the paper does not technically qualify as opinion but probably fall under that category. Thus, only two seemed like real mistakes in the algorithm--a opinion piece about Hilary Clinton and a news article about violence in Brazil.

## Feature selection

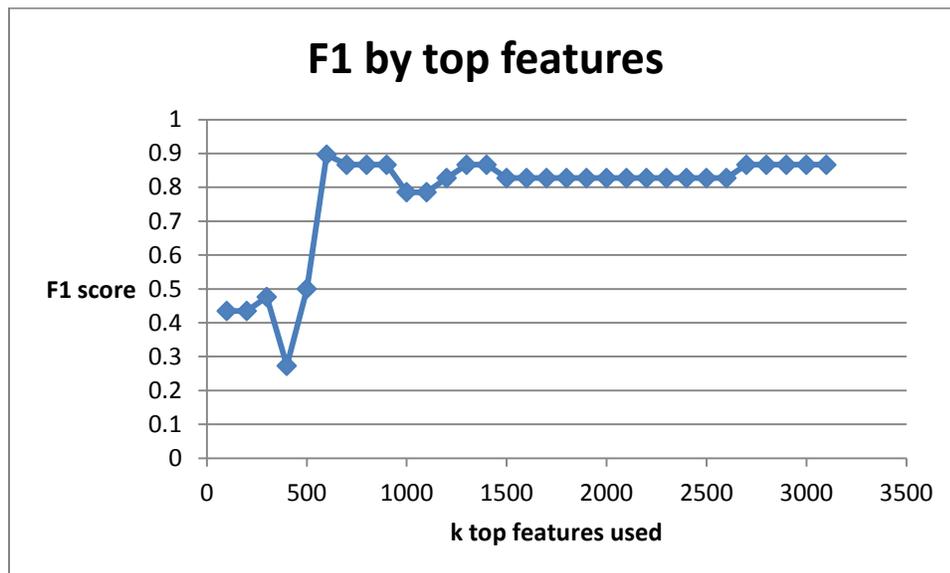


Figure 1: Searching top features (by mutual information score) for optimal number of features

For feature selection, we began with unigram counts using stop-words. With just unigram counts alone, an SVM achieved .85 F1. Sparsity is a common problem with unigram models in which the number of features is much less than the number of training examples ( $N_g$ ). Thus, we tried several successful techniques for reducing the feature space. With porter stemming, the score increased to .88. Using a mutual information measure of binarized features vectors, we searched the space of number of features in increments of 100, peaking at 600 features and an F1 of .9. This indicates that the top 600 words are better for distinguishing opinion from news than the space of all of the features. Realistically, the differences after 600 features are negligible and the classifier performed with little variation for all counts of features tested from 600 onwards.

We also tried using TF-IDF instead of counts. Previous research has suggested that TF-IDF improves scores for text classification (Rennie et al, 2003; Toprak and Gurevych, 2009). We were unable to replicate these results and instead saw F1 decreased to .70. While we do not have a good explanation for why this should be different, usage of stop words and stemming might have helped eliminate words like "the" that would be overcounted. The goal of TF-IDF is to give higher weight to words that occur a lot in a document but little over the corpus. Another theory is that if a news and opinion piece are about the same event, they will have high TF-IDF for words related to that event but that word will not help to distinguish the class. However, the phenomenon is likely a peculiarity of the dataset.

Finally, some work showed that part of speech counts might be effective at subjectivity classification (Toprak and Gurevych, 2009). To test this, we used the counts of part of speech tags from the Penn Treebank tagger. The results were unsuccessful with F1 falling to .67, with articles mostly getting classified as News. Nor did these improve score when used in conjunction with unigram features.

## Language-related results

Top-rated by mutual information for short-term dataset:

- |            |               |
|------------|---------------|
| 1. quot    | 11. tax       |
| 2. year    | 12. officials |
| 3. party   | 13. city      |
| 4. years   | 14. medicaid  |
| 5. israel  | 15. court     |
| 6. federal | 16. women     |
| 7. bbc     | 17. campaign  |
| 8. united  | 18. cuts      |
| 9. ms      | 19. country   |
| 10. time   | 20. american  |

Top-rated by mutual information for long-term dataset:

- |            |                |
|------------|----------------|
| 1. dr      | 11. studi      |
| 2. report  | 12. kill       |
| 3. work    | 13. told       |
| 4. product | 14. republican |
| 5. percent | 15. street     |
| 6. iraq    | 16. research   |
| 7. includ  | 17. plan       |
| 8. world   | 18. polic      |
| 9. project | 19. program    |
| 10. secur  | 20. rais       |

The world that mutual information measurement found to be most informative of category corroborated the hypothesis that traditional sentiment lexicons such as TUD subjective verb lexicon used in Toprak and Gurevych to some would not be as effective for news articles (Toprak and Gurevych, 2009; TUD).

The short term data set as expected includes more words related to specific news events of the last few months--especially politics related ones that were prevalent during the United States election season, such as *campaign*, *country*, *party*, and the word *israel* due to the Israeli attack on Gaza. In the short term, particular news pieces are more successful than opinion or news related words in general at distinguishing the categories.

The long term data set, by contrast, included only one word that appeared to be related to a particular event: *iraq*. Since the Iraq war lasted over the entire period that the dataset was collected from, the presence of the word makes sense. The rest of the words, such as *report*, *work*, *percent*, *kill*, or *polic* seem to be clearly connected reporting or opining.

I had hypothesized that the classifier would do better on the short-term dataset than the long term with the rational that there were a lot more samples for a short time period, so in theory there was more training data. However, the short-term performed at about the same level with an F1 of .87. The short-term dataset would be more prone to the type of event-related issues discussed in relation to TF-IDF. I also note that a multinomial Naive Bayes model performed better on the

short-term and an SVM performed better on the long term. Other results suggest that in general SVM outperforms Naive Bayes in text classification (Rennie, 2003).

### **Conclusions**

The task of distinguishing opinion and news appears to be ones that can be solved with relatively simple tools, much to the credit of the *New York Times*. Prior work in document classification appears to have been effective at this specific classification task. In the future, it would be interesting to explore generalizing the task to different dataset to test whether the lexicon of news/opinion words generated by the model succeeds in classifying articles from other newspapers, news sources, blogs, etc. One could also try using features related to sentence structure. These would be unlikely to improve score but might provide interesting linguistic insights.

### **Sources**

Liu. "Sentiment Analysis and Subjectivity." Handbook of Natural Language Processing, 2nd Edition. 2010.

Ng. CS229 class notes. CS299.stanford.edu.

Pang and Lee, "Opinion mining and sentiment analysis." Foundations and Trends in Information Retrieval 2(1-2), pp. 1–135, 2008.

Potts. "Sentiment Analysis Tutorial." <http://sentiment.christopherpotts.net/>

Rennie et al. "Tackling the Poor Assumptions of Naive Bayes Text Classifiers." 2003.

Sebastiani. "Machine Learning in Automated Text Categorization." ACM computing surveys. 2003.

Toprak and Gurevych. "Document Level Subjectivity Classification Experiments in DEFT'09 Challenge." DEFT'09. 2009.

TUD subject verb lexicon. <http://www.ukp.tu-darmstadt.de/data/sentiment-analysis/subjective-verbs-lexicons>

Turney and Littman. "Measuring Praise and Criticism: Inference of Semantic Orientation from Association." ACM Transactions on Information Systems, Vol. 21, No. 4, October 2003.

Yu and Hatzivassiloglou. Towards answering opinion questions: separating facts from opinions and identifying the polarity of opinion sentences. In Proceedings of the 2003 Conference on Empirical Methods in Natural Language Processing, pages 129–136, Sapporo, Japan. 2003.

### **Libraries used**

sklearn, nltk