

Identification of Fine Grained Feature Based Event and Sentiment Phrases from Business News Stories

Brett Drury
LIAAD-INSEC
Rua de Ceuta, 118, 6^o
Porto, Portugal
bdrury@liaad.up.pt

J.J. Almeida
Escola de Engenharia
University of Minho
Braga, Portugal
jj@di.uminho.pt

ABSTRACT

The analysis of business / financial news has become a popular area of research because of the possibility to infer the future prospects of companies, economies and economic actors in general on information contained in the media. The classical approaches rely upon a "coarse" polarity classification of a news story, however this may not be an optimal solution because this form of classification assigns the same polarity to all of the entities contained in the news story. A news story which contains multiple entities may contain varying polarity for each individual entity. In addition, "coarse" classification may ignore sentiment modifiers which may alter the strength or direction of the story's polarity. News stories don't have a preassigned polarity label, consequently news stories must be manually assigned a polarity label. This process is slow, therefore there will be limited labelled data available. This lack of pre-classified data may inhibit the performance of learners which rely upon labelled data.

This paper describes a rule based approach which identifies feature based sentiment and business event phrases. The phrases are captured with context free grammars which model the phrase as a triple. The triple contains: 1. Phrase subject (an economic actor), 2. A sentiment adjective or event verb and 3. An object (a property of the phrase subject). The captured phrases are limited by the semantic role of the subject. An annotated phrase can capture sentiment modifiers and negators. The scoring of the phrase incorporates all relevant linguistic features and consequently an accurate individual polarity score can be assigned to each relevant entity. The evaluation of the technique reports a recall of 0.71 and precision of 0.94 sentiment phrase annotation and 0.84 recall and 0.83 precision for event phrase annotation.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WIMS'11, May 25-27, 2011 Sogndal, Norway

Copyright © 2011 ACM 978-1-4503-0148-0/11/05... \$10.00.

Keywords

Sentiment Detection, Event Detection, Linguistic Patterns, News, Business, RSS, Web Intelligence

1. INTRODUCTION

The analysis of business news has become an area of great interest to both the academic researcher and the commercial practitioner because of the possibility of inferring future prospects of economic actors based upon information contained in news. A common approach for approximating the influence of news story it is to estimate its polarity: positive or negative. The "total effect" of news story on a "targeted economic actor" is calculated by multiplying the story's relevancy by its polarity category (+1/-1). This approach has produced positive results [12], however it is not an optimum solution because it fails to: 1. Account for the difference between events [9] and sentiment, 2. Differentiate between general sentiment and more specific feature sentiment and 3. Correctly assign sentiment to the target entity when the news story refers to more than one economic actor.

News polarity classification normally relies upon the identification of sentiment / opinion words which are normally adjectives. The two most common approaches to sentiment classification are: machine learning and general sentiment dictionaries for example: SentiwordNet [10]. Machine learning techniques require training data which has been sorted into negative and positive categories. The learner then infers categories for new uncategorised data based upon features learnt from the training data. There are two problems with this approach: 1. Hand collating training data can be a laborious task, 2. There is no guarantee that the training data is an accurate representation of the domain. Semi-supervised learning (SSL) can be used to increase the amount of training data and consequently improve the accuracy of the learner, however in a number of circumstances SSL can decrease the accuracy [1, 6]. General sentiment dictionaries contain pre-scored words. This score can be an indication of the strength of the word's polarity. A combination of scores may provide an indication of the polarity

of a news story. These dictionaries are not domain specific and consequently may underplay or omit domain specific language and in specific circumstances perform poorly [3]. Sentiment dictionaries are not useful for event detection because events lack sentimental /opinionated language.

The described method annotates the news story with: event and sentiment phrases. This information can be used to provide event and sentiment summaries of the news story which the economic literature suggests may be an important factor [9] for inferring the future prospects of "economic actors". These summaries may improve machine learning techniques when there is limited training data because "redundant information" is removed from the both the labelled and unlabelled data. The event and sentiment phrases are described as triples: 1. Economic Actor, 2.Sentiment Adjective or Event Verb, 3. Economic Actor Property. The triples are described as JAPE Grammars. Jape is a regular expression tool for annotations [8]. The paper also describes strategies for phrase annotation when the "Named Entity" element is missing. The technique when tested against a small Gold Standard reports: a recall of 0.71 and a precision of 0.94 for extracting sentiment phrases and a 0.84 recall and 0.83 precision for extracting event phrases. The remainder of the paper will describe the following: News Story Acquisition, Domain Specific Lexicon Construction and Analysis , Grammar Rule Induction and Evaluation.

2. RELATED WORK

A review of the research literature reveals a number of approaches for extracting information from a document collection: templates (regular expressions), supervised learning and semi-supervised learning [20]. A template identifies the types of items which are required to be extracted from the text, for example Person and Position. Templates are manually constructed by human experts which can be time consuming [20]. Supervised learning relies upon a document collection where phrases of interest have been manually annotated. The annotation process can be a time consuming process [20]. The semi-supervised approach to information extraction require less annotated text, but is a harder learning problem [20]. The Autoslog systems [19] extracted phrase patterns for noun phrases in the training corpus, and use co-occurrence to rank the extraction patterns. These systems and others were designed to extract events for specific types of events, for example management transition. The papers did not describe any attempt to score or classify the phrases and did not describe experiments for the extraction of sentiment phrases. The extraction of sentiment phrases is popular in the "product review" domain and is often referred to as featured based sentiment or featured based summaries [17]. A "feature" of a product, for example fuel economy of a car, can have a specific sentiment which may be contrary to the general sentiment for the product. Bing [17] describes a supervised approach - Label Sequential Rules - (LSR) which learns sequences of Part of Speech (POS) tags to extract sentiment phrases for specific features of a given product. It was not possible to locate in the research literature the application of the LSR approach to financial news.

2.1 Research Goals and Contributions

The motivation of this work was to construct a rule based system which: 1. extracted sentiment and event phrases, 2.

locate target of sentiment or event phrase, 3. provide distinct scoring methodologies for event and sentiment phrases, 4. provide a phrase combination strategy when an phrase element is missing. A secondary motivation was to construct the rules in such a manner that there were no restrictions in the types of phrases to be returned. The literature search indicated that previous research had a narrow criteria and consequently these systems would discover phrases of a known type, for example management transition. There was a general interest in "business event and sentiment phrases", but not in a specific type of business phrase.

The contribution of this paper is to provide a simplified rule based method which identifies features at the phrase level which may provide a basis for making inferences about the economic actor. In the review of the literature it was not possible to locate a methodology which differentiates between sentiment and events. The economic literature suggests that the effects of events and sentiment are different; the impact of events is immediate where as sentiment has an accumulative effect over time. The separation of the two types of phrases allows separate inferences to be made concerning an economic actor. A further contribution of this paper was to use a feature based approach with business news rather than the product review domain. The final contribution of this paper is to provide a method of combining partial phrases (when there is an element missing) and the identification of the target of the phrase when the phrase makes no direct reference to it.

3. NEWS STORY ACQUISITION

News stories were gathered from freely available sources published on the Internet. The news story was accessed from information described in Really Simple Syndication (RSS) format on the publisher's site. It was possible to extract the following meta information: headline, description, category information from the RSS file. The story text was extracted from the news story HTML; meta-data for each story was provided by the Open Calais [18] Web Service.

3.1 Removal of Duplicate and Non-Financial News Stories

The corpus contained in-excess of 200,000 news stories; although the stories were gathered from financial RSS feeds there were a large number of stories which were non-financial. A number of stories were duplicated although they had different story url and publication dates. Duplicate stories were removed by the comparison of each story's RSS:headline and RSS:description fields with the existing stories's RSS:headline and RSS:description; if two stories or more had the same headline and description fields then all but one story was removed. A category for each news story was contained in the Open Calais Meta-Data [18]. If the news story was not categorized as financial or business news then it was removed from the training set. The remaining stories will be known as the "Training Stories".

The "Training Stories" text was split into sentences with the ANNIE Sentence Splitter [7]. The following named entities were extracted from the meta-data: companies; organizations, market indexes and company employees. These types

of entities for the purposes will known as financial named entities (FE). A sentence was removed from the training set if it did not contain one of the aforementioned entities. The process of removing sentences reduced the number of training sentences to approximately 500,000; this set of sentences will be known as "Training Sentences". 770 sentences were reserved and annotated with the following information: 1. Event phrase and direction, 2. Sentiment phrase and direction, 3. Sentiment / Event phrase target. This group of sentences will be known as "The Gold Standard".

4. DOMAIN LEXICON CONSTRUCTION AND ANALYSIS

The identification of event phrases was predicated upon the discovery of event verbs; sentiment extraction relied upon the identification of opinionated words which the research literature indicates are normally adjectives [14]. The "Training Sentences" set was parsed with the ANNIE Part of Speech (POS) Tagger [7] to assign part of speech information to each word in the "Training Sentences" set. The POS information was used to identify possible event and sentiment candidate words.

4.1 Extraction of Event Verbs

As previously stated verbs can be an indication of events [16]; consequently Verbs were extracted from the POS tagged "Training Sentences". The Verbs were extracted (base verbs) and sorted by frequency. The verbs which describe the actions of economic actors, for example the verb "rise" (Microsoft's shares will rise), will be known for the purpose of this paper as "event verbs". A subset of the most frequent "event verbs" were hand-selected from the base verbs, this set of verbs will be known as "base event verbs". The base event verbs were expanded with similar verbs from the Levin verb categories [16], for example the Verb "bounce" was part of the Roll Verbs category [16], consequently it was possible to expand "bounce" with the following words: "drift", "drop", "float", "glide", "move", "roll", "slide", "swing". This expanded list will be known as "expanded event verbs". The "expanded event verbs list" was further expanded with semantic equivalents from Wordnet [11] of each set member. This expanded list will be known as "final event verbs".

A domain expert categorized¹ and "scored" the "final event verbs". The "positive" verbs were assigned a symbolic score of "+1" and the "negative" verbs were given a nominal value of "-1". Examples of the verbs and their categories are shown in Table 1.

Verb Category	Examples
Obtained	gain(+), add(+), forge(+), win(+), attract(+)
Lost	fire(-), cut(-), cancel(-)
Direction	climb(+), fall(-), boost(+), down(-)
Behaviour	storm(+), unravel(-)
Influence	hurt(-), hit(-), push(+), suffer(-)

Table 1: Sample Verb Extraction Categorization

¹Categorization was for clarification purposes only

4.2 Extraction of Sentiment Adjectives

Sentiment information can be contained in adjectives [5], for example: 1. worst (negative), 2. good (positive). Adjectives were consequently extracted from the POS tagged sentences. The most frequent positive and negative adjectives were extracted², this list will be known as "seed adjective list". The "seed adjective list" was expanded by extracting semantic equivalents from WordNet [11]. This expanded list was further extended with algorithm 1, which is a simplified version of the algorithm described by Hatzivassiloglou and Columbia [13] which uses connectives to identify and predict new sentiment words and their orientation. Connectives were used in this instance to identify adjectives with the same sentiment label. For example: if the adjective "good" has a known polarity label (+) it is possible to propagate it's label to other words, the following sequence was extracted from our corpus: 1.good and cost-efficient, 2.cost-efficient and fair, 3.fair and transparent. In this sequence the label was propagated from "good" to: 1. cost-effective, 2. fair, 3.transparent with the connective "and".

Algorithm 1: Description of Sentiment Propagation

Input: SL: A list of adjectives with sentiment labels

Input: UD: List of unlabelled sentences

Input: LC: List of Connectives

Output: SL: Expanded list of adjectives

$trigs \leftarrow \text{calctrigrams}(UD)$

repeat

$candidates \leftarrow \{\}$

for $(word, lab) \in SL.words$ **do**

for $\{(w_1, w_2, w_3) \in trigs | w_1 = word\}$ **do**

 next if $w_2 \notin LC$

 next if $w_3 \in SL$

$push(candidates, (w_3, lab))$

$SL \leftarrow SL \cup candidates$

until No more new candidates

Each iteration of the algorithm produced new opinion words which were not in the input list of words; the new words were expanded with semantic equivalents from WordNet. A new input list which consisted of the newly expanded words was created and used as a seed list for the new iteration of the algorithm. This process was continued until no more new words were produced. The positive and negative words were assigned scores of "+1" and "-1".

4.3 Extraction of Entity Features

The initial set of training sentences was reduced by eliminating sentences which did not contain either: an event verb, a sentiment adjective. This new set of sentences will be known as "reduced training set". The "reduced training set" was used to identify words which had a statistical significant relationship with the identified event verbs or sentiment adjectives. The sentences contained in the "reduced training set" had the following word types removed: stop words, proper nouns, named entities. The remaining words were extracted and labelled with one of the following categories:

²The most frequent adjectives were scored as positive or negative with information from Sentiwordnet

co-occurred with event verb, co-occurred with sentiment adjective, co-occurred with both sentiment adjective and event verb. The remaining words will be referred to as the "co-occurring word list".

A Mutual Information (MI) score was calculated for each member of the "co-occurring word list". The following were removed from the "co-occurring word list": 1. Words which MI indicated co-occurred by chance, 2. Event verbs and sentiment adjectives. The collection of the remaining words will be referred to as "statistically significant word list". A new expanded list was created with all the word forms of "statistically significant word list" members. A sample of the members of the "statistically significant word list" can be found in Table 2.

Categorization	Examples
Success Measures	footfall, sales, profits, demand
Time Periods	Monday, Tuesday, January, month, year, period
Third Parties	investors, analysts, investors, economists, regulators, consumers
Miscellaneous	transactions, finance, bankruptcy

Table 2: Sample Features (Nouns)

4.4 Extraction of Sentiment and Event modifiers

The strength of a "sentiment word" may be modified by an adverb [5]. Sentiment modification maybe be one of the following: sentiment maximization (e.g. very good), sentiment minimization (e.g. fairly good) and negation (e.g. not good). The same procedure for identifying co-occurring nouns was used to identify sentiment modifying adverbs. The adverbs were hand-scored, negation words with -1, minimizers with a score of between 0.1 and 0.9 whilst maximizers with a value between 1.1 and 2. Table 3 contains some examples.

Sentiment Categorization	Modifier	Examples
Maximization		sharply, super, perfectly
Minimization		rickety, piffling, just
Negation		not, none, never

Table 3: Sentiment Modification (Adverbs)

4.5 Lexicon Construction Summary

The above process produced the following: 1. 2519 adjectives with a polarity label, 2. 393 verbs with a polarity features, 3. 2609 entity features, 4. 90 sentiment modifiers.

5. GRAMMAR RULE INDUCTION FOR EVENT AND PHRASE ANNOTATION

This paper has thus far described the identification of words which have statistical relations with either an: Event Phrase or Sentiment Phrase. The motivation of this work was to

annotate event or sentiment phrases, consequently it was necessary to construct a series of grammars (rules). The grammars were expressed in JAPE, which can manipulate annotations in GATE.

5.1 Gate Annotations

JAPE Grammars require annotations to manipulate. Words can be annotated in GATE with the GATE Gazetteer, the GATE Gazetteer holds a list of pairs: a word and it's annotation label. The Gazetteer was supplemented with the following: adverbs (sentiment modifiers), event verbs (verb); sentiment words (adjective) and statistically significant words (FE features) as new word lists. The Gazetteer already contained a company list which was expanded with FEs from the Open Calais meta-data. The elements of event and sentiment phrases were now annotated in the story text and could be manipulated by JAPE.

5.2 Phrase Extraction Patterns

The JAPE grammars were based upon extraction triples. The JAPE rules were not order dependent, for example the event extraction pattern could be: Named Entity, Verb, Feature (Microsoft(NE) reported a drop(Verb) in profits(Feature)) or Named Entity, Feature, Verb (Microsoft's(NE) profits(Feature) dropped(Verb)). The grammars ensured that the longest possible phrase was returned.

5.2.1 Third Party Financial Entities

There were two types of FEs: 1. The subject of the phrase or 2. A third party passing commenting. Named entities which have a linguistic cue that they were a third party, for example "Analysis by Lane Clark & Peacock, the actuarial consultants" [22], were excluded from analysis.

5.3 Partial Patterns and Backtracking

A number of event and sentiment phrases in the corpus did not contain a FE (company, market index etc). The title of the named entity would either be: implied, substituted with an informal name. Two strategies were followed to compensate for a missing FE: backtracking and partial patterns.

5.3.1 Backtracking

Frequently journalists replace company names in text which makes numerous references to the same named entity (company, market indexes etc) to ensure that the text is not repetitive. The backtracking strategy looked for an explicit reference to a named entity, which was not a third party passing comment, in the previous sentence. The annotated phrase was not expanded, but the inferred named entity was used as the event or sentiment target. In certain circumstances the partial phrase was expanded with partial pattern combination.

5.3.2 Partial Patterns

When the FE element was missing from the immediate sentence, the remaining elements (Verb and Feature or Adjective and Feature) were returned. The complete event and sentiment phrase was returned by combining two or more partial patterns in the same sentence. There were two combination rules:

Rule 1 - Partial patterns were joined when there was one separator token (space; carriage return, new line etc) between partial patterns.
Rule 2 - Partial patterns were joined when they were separated by a "continuation" [4].
Note - The patterns must the same "type", event phrases can't be joined with sentiment phrases.

Table 4: Combination Rules

5.4 Annotation of Complete Phrases

Thus far there have been two types of patterns described: complete patterns and partial patterns. Complete patterns and combinations of partial patterns may capture whole event or sentiment phrases. On occasions these patterns may not be sufficient, consequently it was necessary to combine patterns, both complete and combined partial patterns. The combination rules were the same as for partial patterns as described in Table 4 .

5.4.1 Examples of pattern combination

Examples of pattern combination are illustrated in Figures 1 and 2. The example in Figure 1 is a typical event phrase which was captured by a combination of complete pattern and a partial pattern. The complete pattern contains a named entity, Lockheed Martin, an event verb "drop" and a series of features. The partial pattern contains a named entity, "S&P500" (which is a name of a financial index) and a comparison word, "outperformed". The two patterns were separated by the continuation "although", therefore the two patterns were joined to create a single phrase.

The sentiment phrase was captured by a single complete pattern; the pattern contained a Subject (a Company Name, "Megitt"), two sentiment words ("exposed" and "hit") and a Feature (period). The extracted phrase contained an event partial pattern, "down over 60% over the period". This was a frequent occurrence, sentiment language often commented on a described event.

5.4.2 Estimating the Polarity of a Phrase

This task was separated into two: event polarity and sentiment direction. Nominally event polarity was estimated by assigning the polarity of the event verb (which was assigned by a domain expert) to the whole phrase, however a number of features (nouns) inverted the score of the phrase, for example the phrase: "A drop in profits" would be negative, but the phrase "A drop in costs" would be positive. Sentiment polarity was achieved through the application of the AVAC algorithm [21]. The AVAC algorithm attempts to use the sentiment modifiers and negation to estimate the sentiment direction of the sentiment phrase.

6. EVALUATION

The evaluation criteria was two-fold: 1. comparison with a manually annotated gold standard, 2. training document classifiers. The "gold standard" evaluation was to determine: 1. precision and 2. recall of the rule based system with an expert's annotations. The second evaluation is a comparative evaluation of the proposed technique with market alignment [15] to train a document classifier. The training of new story classifier was a secondary aim of this technique.

6.1 Gold Standard Evaluation

The evaluation of the phrases annotated by the grammars was by comparison with the previously mentioned "The Gold Standard" set of sentences. The evaluation consisted of the following tasks: 1. Correct identification of the sentiment or event phrase, 2. Differentiation of an event from sentiment, 3. Correct identification of sentiment / event target and 4. Direction of sentiment or event. The rational of the evaluation criteria is described in Table 5.

Phrase Extraction	Extracted Phrase Must Convey The Fundamental Message of The Annotated Phrase
Event / Sentiment Detection	Must Determine Between A Sentiment or Event Phrase. If Annotation is Uncertain then Either Will be Accepted
Target of Sentiment	Extracted Sentiment Target Must Exactly Match The Annotation
Sentiment / Event Direction	Extracted Must Match The Direction Of The Annotation

Table 5: Evaluation Criteria

Evaluation Item	Recall	Precision
Sentiment Phrase Extraction and Direction	0.71	0.94
Event Phrase Extraction and Direction	0.84	0.83
Sentiment Target Extraction	0.74	0.74
Event Target Extraction	0.84	0.77

Table 6: Recall and Precision for Phrase Extraction

6.1.1 Annotation

There was a single annotator. The annotator annotated each sentiment and event phrase with a polarity and a target, for example:

<start>A recent analysis of the UK industry by Numis Securities makes clear that both commercial and defence companies are well financed, with little or no debt rolling over<end> <Sentiment=positive target=UK Industry> [22].

The annotations are in angled brackets. The example demonstrates: 1. The start and end of a phrase, 2. The type of phrase, 3. The direction of the phrase and 4. The target of the phrase. The annotation process was slow due because: 1. The sentences could be detailed and long, 2. There was no annotation tool and the annotation were typed by the annotator.

6.1.2 Results

The results of the evaluation are presented in Figure 6. The assigned direction of the extracted sentiment phrase was correct for 77% of the extracted phrases, and the assigned direction of the extracted event phrase was correct for 86% of the extracted phrases. It should be noted that for some phrases it was difficult for the annotator to clearly define an event from sentiment.




Figure 1: Annotated Event Sentence [22]




Figure 2: Annotated Sentiment Sentence [22]

A manual approach was also evaluated. The manual approach used regular expressions to capture words in certain part of speech tags (POS) sequences. This approach returned a negligible number of phrases because the manually identified sequences were a very limited subset of the total number event and sentiment POS sequences contained in the corpus.

6.1.3 Supplementary Results

Another set of documents which had been separately manually annotated for alternative experiments with sentiment information at sentence level were available for evaluation. The document set contained texts which were from the business domain, but were more general in nature as it included macro-economic news and government announcements. The application reported an accuracy of 0.77 for extraction and scoring of sentiment phrases. When sentences which had no sentiment information and consequently no phrase was extracted were included in the calculation the figure rose to 0.85.

The evaluation procedure revealed that on average there were 22 sentiment or event phrases extracted per news story. If the evaluation documents were representative of the corpus then there are a possible 440000 sentiment or event phrases contained in our text collection.

6.2 Comparative Document Classification

The comparative document classification evaluation task was designed to evaluate two identical learners whose training

set of documents were generated by two separate methods: 1. market alignment, 2. the proposed rule based method. The FTSE 100³ was chosen as the aligning market. News stories which were published in a U.K. newspaper and were categorized as UK or British business news and FTSE 100 market data was gathered from free sources on the Internet from January 2009 until June 2010. Stories which coincided with a market movement of 100 points or more points were assumed to be positive and a market drop of 98 or more points was assumed to be negative. The proposed method classified stories by evaluating their headline. Headlines provide a good indication [2] of the following news story, consequently a news story with a positive headline will have positive story text. Headline was considered to be positive or negative if the proposed method returned: 1. positive or negative event phrase, 2. positive or negative sentiment phrase. The models induced by the two learners were evaluated with a 2 * 5 cross fold validation. The results are presented in table 7.

The performance of the market alignment approach was poor; this result was unsurprising as news stories published on the web are not "timely", i.e. "positive" news stories can be published on days that the market has lost points and "negative" news stories can be published on days the market has gained points. The proposed method is consistent in its selection of positive and negative stories, and consequently the model induced by the learner from data selected by the

³More information concerning the index can be found here:<http://goo.gl/xUwh>

Method	F-Measure	Recall	Precision
Market Alignment	0.57 +/- (0.01)	0.57 +/- (0.01)	0.57 +/- (0.01)
Proposed Method	0.77 +/- (0.01)	0.76 +/- (0.01)	0.77 +/- (0.01)

Table 7: Comparative Evaluation of News Story Classifiers

proposed method is more robust.

7. CONCLUSION

Financial news unlike a other forms of written language is published to inform rather than entertain. The constraints of producing clear and informative text which conforms to standard grammatical conventions removes a number of obstacles for computational analysis of the news story. The financial news domain is a restricted domain and therefore although the economic actors in the news stories may change, the general themes are similar. The restricted nature of this domain coupled with relatively unambiguous and clear text may allow simple techniques to produce relatively good results. It is likely that more sophisticated techniques may produce better results, however this rule based approach may provide a basis for further investigation of news stories.

7.1 Future Work

The initial focus of the future work will be to expand the gold standard. The current gold standard is small because the annotation process is a manually intensive task. The phrase extractor will be used in future research which will include: 1. Training alternative classifiers, 2. Sentiment / event based information recall and 3. Event / Sentiment summaries for news stories.

8. REFERENCES

- [1] Steven Abney. *Semisupervised Learning for Computational Linguistics*. 2007.
- [2] Blake C. Andrew. Media-generated shortcuts: Do newspaper headlines present another roadblock for low-information rationality? *The Harvard International Journal of Press/Politics*, 12(2):24–43, 2007.
- [3] Alexandra Balahur, Ralf Steinberger, Mijail Kabadjov, Vanni Zavarella, Erik van der Goot, Matina Halkia, Bruno Pouliquen, and Jenya Belyaeva. Sentiment analysis in the news. In Nicoletta Calzolari, editor, *Proceedings of the Seventh conference on International Language Resources and Evaluation*, Valletta, Malta, may 2010.
- [4] Chris Barker. Continuations in natural language. In *Fourth ACM SIGPLAN Continuations Workshop (CW'04)*.
- [5] Farah Benamara, Carmine Cesarano, Antonio Picariello, Diego Reforgiato, and V. S. Subrahmanian. Sentiment analysis: Adjectives and adverbs are better than adjectives alone. In *Proceedings of the International Conference on Weblogs and Social Media (ICWSM)*, 2007.
- [6] O. Chapelle, B. Schölkopf, and A. Zien, editors. *Semi-Supervised Learning*. MIT Press, Cambridge, MA, 2006.
- [7] Bontcheva Cunningham, Maynard and Tablan. Gate: A framework and graphical development environment for robust nlp tools and applications. In *Proceedings of the 40th Anniversary Meeting of the Association for Computational Linguistics*, 2002.
- [8] H. Cunningham, D. Maynard, K. Bontcheva, and V. Tablan. Gate: A framework and graphical development environment for robust nlp tools and applications. In *Proceedings of the 40th Anniversary Meeting of the Association for Computational Linguistics*, 2002.
- [9] Werner F. M. De Bondt and Thaler R. Does the stock market overreact? *The Journal of Finance*, 40(3):793–805, 1985.
- [10] A Esuli and F Sebastiani. Sentiwordnet a publicly available lexical resource for opinion mining, 2006.
- [11] C. Fellbaum. *WordNet: An Electronical Lexical Database*. The MIT Press, Cambridge, MA, 1998.
- [12] G. Gidofalvi. and C. Elkan. Using news articles to predict stock price movements. Technical report, University of California, 2003.
- [13] Vasileios Hatzivassiloglou and Kathleen R. McKeown. Predicting the semantic orientation of adjectives. In *Proceedings of the eighth conference on European chapter of the Association for Computational Linguistics*, pages 174–181, 1997.
- [14] Nitin Indurkha and Fred J. Damerau. *Handbook of Natural Language Processing*. Chapman & Hall/CRC, 2010.
- [15] Victor Lavrenko, Matt Schmill, Dawn Lawrie, Paul Ogilvie, David Jensen, and James Allan. Language models for financial news recommendation. In *Proceedings of the ninth international conference on Information and knowledge management, CIKM '00*, pages 389–396, New York, NY, USA, 2000. ACM.
- [16] Beth Levin. *English verb classes and alternations: a preliminary investigation* by Beth Levin. The University of Chicago Press, 1993.
- [17] Bing Liu. *Handbook of Natural Language Processing*. Springer, 2007.
- [18] Reuters. Calais web service. <http://opencalais.com/>, consulted in 2009.
- [19] Ellen Riloff and Jay Shoen. Automatically acquiring conceptual patterns without an annotated corpus. In *In Proceedings of the Third Workshop on Very Large Corpora*, pages 148–161, 1995.
- [20] Mark Stevenson and Roman Yangarber.
- [21] V. S. Subrahmanian and Diego Reforgiato. Ava: Adjective-verb-adverb combinations for sentiment analysis. *IEEE Intelligent Systems*, 23(4):43–50, 2008.
- [22] Financial Times. Why defence should prove to be defensive. <http://www.ft.com/cms/s/0/91e11db0-9ee1-11dd-98bd-000077b07658.html>, consulted in 2009.