



HAL
open science

A Contextual Classification Strategy for Polarity Classification of Direct Quotations from Financial News

Brett Drury, Gaël Dias, Luis Torgo

► **To cite this version:**

Brett Drury, Gaël Dias, Luis Torgo. A Contextual Classification Strategy for Polarity Classification of Direct Quotations from Financial News. International Conference On Recent Advances in Natural Language Processing (RANLP 2011), Sep 2011, Hissar, Bulgaria. pp434-440. hal-01069030

HAL Id: hal-01069030

<https://hal.science/hal-01069030>

Submitted on 1 Oct 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Contextual Classification Strategy for Polarity Analysis of Direct Quotations from Financial News

Brett Drury
LIAAD-INESC
Portugal
brett.drury@gmail.com

Gaël Dias
HULTIG, Portugal
DLU/GREYC, France
ddg@di.ubi.pt

Luís Torgo
Fac. Sciences
LIAAD-INESC, Portugal
ltorgo@inescporto.pt

Abstract

Quotations from financial leaders can have significant influence upon the immediate prospects of economic actors. Indiscreet or candid comments from senior business leaders have had detrimental effects upon their organizations. Established polarity classification techniques perform poorly when classifying quotations because they display a number of complex linguistic features and lack of training data. The proposed strategy segments the quotations by inferred “opinion maker” role and then applies individual polarity classification strategies to each group of the segmented quotations. This strategy demonstrates a clear advantage over applying classical classification techniques to the whole corpus of quotations. While modelling contextual information with Random Forests based on a vector of unigrams plus the “opinion maker role” reaches a maximum F-measure of 52.85%, understanding the “bias” of the quotation maker previously based on its lexical usage allows 86.23% F-measure for “unbiased” quotations and 71.10% F-measure for “biased” quotations with the Naive Bayes classifier.

1 Introduction

Quotations from business leaders or government ministers can have profound effects upon the immediate and future prospects of economic actors. This phenomenon was demonstrated in a 1991 speech by Gerald Ratner at the Institute of Directors. He described his company’s products as “crap” (Ratner, 2007) and that a pair of earrings sold by his company were “cheaper than a prawn sandwich but probably wouldn’t last as long” (Ratner, 2007). His company (Ratners) lost

500 million pounds in value and had to change its name to Signet to distance itself from his speech. There are other, less colourful, examples of quotations impacting the financial prospects of an economic actor. Mervyn King, the governor of the Bank of England, declared in 2008 that “*now seems likely that Britain is entering a recession*”¹. The day after, the British Pound promptly lost value on the foreign exchange markets.

Quotations are arguably an important source of information for researchers trying to determine the financial prospects of an economic actor. However, analysing quotations in terms of conveyed opinion is a non-trivial task because (1) quotations may contain metaphors, euphemisms, slang, obscenities, invented words or negations and (2) their polarity mainly depends on the context of the quotation and in particular, the opinion maker.

Most of the strategies proposed so far for polarity classification or opinion mining have been focusing on positively or negatively labelling word/phrases (Hatzivassiloglou and McKeown, 1997; Strapparava and Mihalcea, 2008), sentences (Hatzivassiloglou and Wiebe, 2000; Turney, 2002) or texts (Pang et al., 2002; Chesley et al., 2006) independent of their context.

Recently, the context of opinion has been addressed and research literature has revealed two different approaches. The first approach was proposed by (Al Masum Shaikh et al., 2007). The underlying idea is to either group the affective information into sets of emotions or to associate the affective information with the opinion of its readers. The contextual information is the reader in contrast to the writer. For example, the following neutral statement in terms of the writer “*Real Madrid won the Spanish Football Cup against FC Barcelona*” can be interpreted as a negative emotion for a Barcelona fan and as a positive one for

¹<http://news.bbc.co.uk/2/hi/business/7682723.stm>, consulted in 2011 (Gloomy forecasts for UK economy).

a Real Madrid fan. These studies show a user-centric approach based on personalization. Second, some works have been emerging, which focus on polarity detection of texts based on contextual information such as the author, the reader or the text. In particular, (Balahur and Steinberger, 2009) identified the main tasks for news opinion mining: (1) the definition of the target, (2) the separation of the good and bad news content from the good and bad sentiment expressed on the target and (3) the analysis of clearly marked opinion that is expressed explicitly, not needing interpretation or the use of world knowledge. In particular, they show that it is important to distinguish three different possible views on newspaper articles: author, reader and text. These have to be addressed differently at the time of analysing sentiment, especially the case of author intention and reader interpretation, where specific profiles must be defined if the proper sentiment is to be extracted. Moreover, (Balahur et al., 2010) presented a work on mining opinions about entities in English language news, in which they tested the relative suitability of various sentiment dictionaries and attempted to separate positive or negative opinions from good or bad news. In their experiments, they tested whether or not subject domain-defining vocabulary should be ignored and results showed that in the context of news opinion mining, subject oriented classification produces better performance than classical strategies.

This paper is concerned with a “market view” of the sentiment in quotations made by an economic actor. A “market view” is expressed either in the rise or fall of a financial instrument or significant increase in trading volume. Classical sentiment classification may assist, but the motivation of the quote maker may inhibit the effectiveness of these techniques as shown in (Balahur and Steinberger, 2009). For example, business leaders lie and when they lie they use opinionated language (Larcker and Zakolyukina, 2010). This characteristic of direct speech will inhibit classical techniques to identify “actionable” information to use in a trading strategy. The research problem is to identify “actionable” information in quotations from financial news.

The proposed approach is predicated upon the following assumptions: (1) certain economic actors are compelled to speak in a highly rhetorical manner which conveys no actionable informa-

tion, (2) rhetorical language contains overtly positive lexicon and (3) certain economic actors are compelled to speak in an objective manner. The final assumption is that an implied role i.e. biased (rhetorical features) or unbiased (non-rhetorical features) can be assigned through job role or specific lexicon extraction.

The proposed approach seeks to group opinion makers by their implied role and apply separate classification strategies to their quotations. This strategy demonstrates a clear advantage over applying classical classification techniques to the whole corpus of quotations. While modelling contextual information with Random Forests based on a vector of unigrams plus the “opinion maker role” reaches a maximum F-measure of 52.85%, understanding the “bias” of the quotation maker previously based on its lexical usage allows 86.23% F-measure for “unbiased” quotes and 71.10% F-measure for “biased” quotes with Naive Bayes.

2 One-Step Learning Strategy

This section will cover the initial experiments and lay some foundations for the justification of the work contained in this paper. The sub-sections will cover the data acquisition process, the learner selection and the influence of the “opinion maker role” of the writer as a feature.

2.1 Data Acquisition

A large number of news stories (>300,000) were collected from freely available sources on the Internet. The news stories were gathered from Really Simple Syndication (RSS) feeds during the period from October 2008 until June 2010. News story meta-data was added by the Open Calais web service². Open Calais identifies quotations, the quotation maker and on occasion job titles and organization affiliations. This process yielded 180,956 quotations, a subset of which were hand-labelled as positive, negative and neutral. The annotation process was conducted by a single annotator. Some examples are given in sentences (1), (2) and (3).

- (1) *Mr Cowgill said the relative strength was a result of the differences between male and female consumers. (neutral)*
- (2) *BBT is trading up on the news as they would likely be able to assume the deposits at an*

²<http://www.opencalais.com/>, consulted in 2011.

attractive price. (positive)

- (3) *About 60 per cent of summer crops could be hurt badly by insufficient rainfall subsequently dragging down agricultural performance which has already been modest in recent quarters.* (negative)

2.2 Baseline Experiments

The assumption of this work is that the role of the quotation maker influences the polarity and the influence over the financial market. The baseline experiments are designed to demonstrate that the addition of the "opinion maker role" as a feature provides a demonstrable gain in F-measure for a classifier. We conducted three different experiments with different feature sets: (1) unigrams, (2) unigrams plus the job role (JR) and (3) unigrams plus the "opinion maker role" (OMR). In particular, the job role was extracted from the Open Calais metadata and the annotator added the "opinion maker role". The annotator selected the "opinion maker role" on the following definitions: (1) *biased* if the opinion maker has a clear affiliation to a company (CEO, CIO etc.) and (2) *unbiased* if the opinion maker is independent of an economic actor and should be free from bias (analysts, economists etc.). The experiments were conducted with the first ranked learner (Random Forests) and a mid ranked classifier (Naive Bayes) based on the classifier ranking assigned by the landmarking tools (Pfahring et al., 2000) implemented in Rapidminer (Mierswa et al., 2006) to ensure that any gain would not be learner specific. The estimated F-measures were calculated with a 10-fold cross validation technique and are presented in Table 1.

Classifier	Features	F-Measure
Rand For.	Unigrams	46.91% \pm 4.07
Rand For.	Unigrams + JR	46.37% \pm 3.06
Rand For.	Unigrams + OMR	52.85% \pm 3.40
N. Bayes	Unigrams	49.01% \pm 4.75
N. Bayes	Unigrams + JR	49.66% \pm 5.10
N. Bayes	Unigrams + OMR	50.54% \pm 4.39

Table 1: Experiments Estimated F-Measures.

The experiments demonstrate a small gain by using the inferred role of the opinion maker, but the gains are within the margin of error. There are, however, gains for both learners and therefore provide some evidence for the "inferred role" assisting the learner. As a consequence, we propose

in the next section the analysis of the language of "biased" and "unbiased" quotation makers in order to see if the "inferred role" can automatically be identified based on a specific language usage and then propose a two-step learning process to improve the accuracy of our learning process.

3 Quote Maker Language Analysis

This section describes the lexical analysis of two job roles: CEOs and Analysts. These two job roles conform to the annotation rules when labelling the baseline experimental data with "opinion maker roles". CEOs are assumed to be part of the "biased" class because they have a direct affiliation with a company whereas analysts are normally independent and therefore are part of the "unbiased" class. If the initial assumption is correct, then the CEOs' quotes would be likely to have rhetorical features whereas the analysts' ones would not.

3.1 The CEOs Lexicon

The expected lexicon of CEOs should contain overtly positive language, which is designed to manipulate the public opinion. The initial lexicon analysis was aimed at extracting adjectives, as adjectives are known to be the conveyors of the opinionated language (Wiebe et al., 2004). For that purpose, we used the Pointwise Mutual Information (PMI) to calculate the affinity of an adjective to a quotation by a person with the job role of CEO. The PMI is defined in Equation 1 where "adj" represents an adjective and "cl" is the job role of CEO.

$$PMI(adj, cl) = \log_2 \frac{Pr(adj, cl)}{Pr(adj)Pr(cl)}. \quad (1)$$

All the adjectives, which scored above zero were assumed to be a member of the CEO's lexicon. As such, 1,401 adjectives were extracted and ranked in order of their PMI score. The majority of the adjectives are positive and there are few negative adjectives. In particular, the first negative adjective is ranked 87. Conversely, the negative language was not exaggerated, however the positive language was domain specific as for example, "win-win", "mission-critical" and exaggerated, as for example, "superb", "immense". In particular, negative adjectives tended to have the lowest PMI scores. A further analysis was made of frequent and infrequent unigrams and bigrams. The analysis was limited to the top and bottom 100 terms.

The most frequent terms were positive whereas the infrequent terms were either atypical words or negative words. In summary, the lexicon of the CEO is overwhelmingly positive, which is, nevertheless contradictory as the quotes were harvested between 2008 and 2010, which was a time of a severe economic crisis.

3.2 Analysts Lexicon

Compared to the CEOs’ language, the Analysts’ language should be more measured because the analysts’ job function is to provide objective advice. The lexicon analysis was the same as for the CEOs i.e. analysis of specific adjectives using the PMI and analysis of frequent and infrequent terms i.e. unigrams and bigrams. The adjective analysis revealed a smaller lexicon, 415 adjectives compared with the 1,401 in the CEO lexicon. The next difference is the higher ranking of negative words. The highest ranking of a negative word is for the adjective “*speculative*”, which had the rank of 2. Comparatively, the highest ranked negative word in the CEO lexicon had a rank of 87. The analysis of frequent and infrequent terms revealed a lack of opinionated language. This is contrary to the CEOs’ language who seems to use positively opinionated language.

In summary, there is clear evidence that there are significant differences in the lexicons of CEOs and Analysts. The lexicon difference in conjunction with the baseline experiments provides justification for the two-step strategy as it will be possible to identify the role of the quotation maker by his language. The next section will describe a methodology to identify “biased” from “unbiased” quotation makers based on their language.

4 Market View of Quotations

The identification of quotations, which contain “actionable information” is a non-trivial task. Manual selection of data is a labourious task and can be impractical because of the volume of information. For example, our data set contained 180,956 quotations. A specific aim of this paper was to identify “real-world” effects of quotations. Consequently, the first attempt to label quotes was to align quotes with market movements as in (Lavrenko et al., 2000). A baseline experiment was conducted where the ticker symbol of the affiliation of the quote maker was retrieved from Yahoo! Finance and the opening and clos-

ing price was recorded. The category of the quote would then be inferred based upon the following conditions: (1) a negative category would be inferred if the share price fell by more than 1%, (2) a positive category would be inferred if the share price rose by more than 1% and (3) a neutral category would be inferred if the share price rose or fell by less or equal than 1%. The evaluation of the automatic labelling was based on a 10-fold cross validation process with the unigrams of the quotations as the only features compared to the manual labelling initially performed. The results are presented in the Table 2.

Learner	Categories	F-Measure
Rand For.	Neut & No-Neut	39.58% \pm 0.0
Rand For.	Neut, Pos & Neg	22.50% \pm 0.0
N. Bayes	Neut & No-Neut	67.42% \pm 4.28
N. Bayes	Neut, Pos & Neg	54.12% \pm 3.70

Table 2: Automatic Labelling F-Measure

The results are clear. Automatic alignment with the market has its flaws. Quotations may appear with a market movement by chance and consequently the inferred label may be false. In fact, this result reproduces other experiments with automatic alignment (Drury et al., 2011). To avoid this kind of problems and achieve acceptable results, auto-alignment of texts with markets requires a form of constraint (Drury et al., 2011). In this paper, we propose a label propagation algorithm for quotations made by an identifiable CEO to improve the automatic labelling of quotations based on the market movement.

4.1 Labelling and Learning CEOs Quotes

The first assumption is that the majority of quotations made by CEOs are likely to be “bluster” and therefore may contain rhetorical language (which is not informative), whereas a small subset would contain useful information. The imbalance between the two categories would ensure that some quotations would “move the market” simply because they would be unexpected. There is some evidence in the research literature to suggest that the element of surprise can move markets (Bomfim and N., 2000). However, surprise is usually infrequent. To confirm our assumption, a human expert aligned a selection of CEO quotations with the movements in the market. The rules for the manual market alignment were the ones explained

above with one difference that the human annotator would make the final decision if the quotation was responsible for the market movement or not. The human annotator found that for every “useful” quotation, there were 100 “bluster” quotations. In fact, it was not possible for the human annotator to align all the CEO quotations in a reasonable period of time. Therefore, once the human annotator had selected a sufficient number of quotations, a further automated process was required. A form of semi-supervised learning was chosen where labels of known data are propagated to unlabelled data via clustering algorithms. The RapidMiner Top Down Clustering operator was chosen as the number of clusters was selected by the operator. The process goes as follows. The seed set of manually annotated quotations is clustered with unlabelled data in groups of 1,000 documents. Clusters with more than 75% of labelled data from a single category have their labels propagated to the quotes in the cluster without labels. This process continues until no further labels are propagated.

This clustering process was executed in three steps: (1) for the initial CEO data (positive), (2) for the quotations attributed to a person with an identifiable job role which was not a CEO (negative) and (3) for quotes attributed to a person with no identifiable job role (neutral). At the end of the process, there were 1,242 quotations which were determined to contain “useful” information. These quotations were split into positive and negative categories with manual alignment with the market. This process was then evaluated based on a 10-fold cross validation with the unigrams of the quotations as the only features showing that regularities can be found as presented in Table 3.

Learner	Categories	F-Measure
Rand For.	Neut & No-Neut	88.28% \pm 2.29
Rand For.	Neut, Pos & Neg	67.10% \pm 2.82
N. Bayes	Neut & No-Neut	82.75% \pm 3.31
N. Bayes	Neut, Pos & Neg	70.71% \pm 3.31

Table 3: Automatic CEO Labelling F-Measure

4.2 Labelling and Learning Analysts Quotes

The role of the analyst has an arguably different role to that of a CEO. Analysts are not required to “bluster” or mislead, and often they tend to reach a consensus (Tamura and Hiromichi, 2003). An analyst consensus ensures that there is a “lack of

surprise” and consequently, a quotation from an analyst is unlikely to move the market. In these conditions, auto-alignment with the market is unlikely to be a profitable strategy. The proposed strategy was to manually extract adjectives from the Analyst lexicon and expand them with WordNet (Miller, 1990) based on existing semantic relationships. The polarity of the adjectives was then inferred by calculating the PMI score for the adjective and its category (i.e. positive or negative) as in (Turney, 2002). As a consequence, in order to collect as strong as possible quotations, a high precision rule classifier selected quotations with three or more adjectives from one category. The classification task was only into positive and negative categories because analysts are assumed not to “bluster” and that the economics of the news publishing business will ensure that quotations will be sufficiently interesting to the general reader before it is published (McManus, 1988). In this case, as there exist many quotations from real-world texts, label propagation was not necessary. Results are shown in Table 4 performed over a 10-fold cross-validation strategy with the unigrams of the quotations as the only features and show how regularities can be found this way.

Learner	Categories	F-Measure
Rand For.	Pos & Neg	83.24% \pm 2.85
N. Bayes	Pos & Neg	86.23% \pm 2.27

Table 4: Automatic Analyst Labelling F-Measure

5 Two-Step Learning Strategy

The initial assumption was that understanding the job role of the opinion maker is likely to lead to improved classification performance upon the contribution of the quotation over the market. On one hand, the quotations with a high level of rhetorical features are assumed to carry no useful information with respect to the financial market. In particular, the quote makers who use rhetorical language are assumed to have the inferred role of “biased” and the groups of individuals who do not use rhetorical language are assumed to be “unbiased”. This was verified in section 3. On the other hand, we know that “biased” people are likely to have loyalties to companies or organizations, whereas “unbiased” people are usually independent because they are employed by companies who provide impartial advice to client. As a con-

sequence, our two-step strategy aims at first learning the “inferred role” of the opinion maker and second applying a unigram classification model to extract positive and negative quotations within the context of the market.

As we showed in section 4, if we are capable of clearly identifying the “inferred role” of the opinion maker, it is likely that we obtain improved performance over classification of quotations as positive, negative or neutral within the context of the financial market. In fact, as shown in section 3, as “biased” and “unbiased” opinion makers use different languages and different linguistic features, learning job roles should be possible. For that purpose, we automatically built a suitable data set through the same clustering process as was used for identifying data for CEOs i.e. the label propagation algorithm. In particular, the data was split into two categories, “bluster”, (i.e. the meaningless category from the CEO data) and “non-bluster”, (i.e. the remaining data from both the analysts and CEO data sets). As a consequence, the clusters, which contained 75% of a single category had their labels propagated and the job titles from the propagated and labelled data were recorded. To better understand the kind of job roles associated to both classes “bluster” and “non-bluster”, we calculated a PMI score for each job title and its affinity to each category. A sample of job titles and their categories are presented in Table 5 and evidence how job role can easily be discovered.

Bluster	Non-Bluster
Chairman, CTO, Co-head, Company President	Chief Economist, Credit Analyst, Chief economic adviser

Table 5: Automatic Identification of Job Roles

The initial assumption is based on the fact that separate strategies take advantage of the individual linguistic characteristics of the hypothesized “inferred roles” in the corpus of the quotation makers. The hypothesized “inferred roles” are in fact “biased” (i.e. quotes made by people with known loyalties to companies/organizations) and “unbiased” (i.e. quotes made by other people without links to companies/organizations). In fact, the market view technique identifies meaningful quotes from the “biased”, but fails to identify quotations from the “unbiased” group because the later group often fails to move the market with their pronounce-

ments. A rule approach works well with the “unbiased” group, but performs worst with the “biased” group because the quotations in the training set are overly positive due to the inclusion of quotes from the “bluster” group. As a consequence, it is compulsory to first identify the “inferred role” of the quotation maker so that the genre specific learner is correctly applied to the given quotation. In fact, the “inferred roles” are based upon job title. The job titles, which have a predominance of rhetorical language and therefore cluster together are for our purposes “biased”. The roles, which have a lack of rhetorical language also cluster together and are assumed to be “unbiased”. So, by applying this two-step strategy, we obtain improved results over the baseline presented in Table 1. In particular, we performed a 10-fold cross validation with the unigrams of the quotations from each category individually as the only features. The results are presented in the Table 6.

Group	Learner	F-Measure
Unbiased	Rand For.	83.24% \pm (2.85)
Unbiased	N. Bayes	86.23 % \pm (2.27)
Biased	Rand For.	64.03% \pm (2.58)
Biased	N. Bayes	71.10% \pm (6.45)

Table 6: Comparison of Inferred Roles

Clustering is a computational expensive process, consequently when classifying a large groups of quotations it is not possible to use this process to separate the quotes into their respective latent groups. The group separation is done by job title as discovered previously. It was therefore possible to accurately separate the potential quotes by keywords into their latent groups. While modelling contextual information with Random Forests based on a vector of unigrams plus the “inferred role” reaches a maximum F-measure of 52.85%, understanding the “bias” of the quotation maker previously based on his job role allows 86.23% F-measure for “unbiased” authors and 71.10% F-measure for “biased” authors with the Naive Bayes classifier.

6 Conclusions

This paper has provided some evidence that grouping quote makers by their latent roles can assist in polarity classification tasks. The paper demonstrates that quote makers latent role pre-determines their language in direct quotations and

consequently quotations by members of these latent roles are susceptible to different forms of analysis. In this paper, we provided evidence of the existence of two latent groups, but we are not arguing that there are only two latent groups in a quotation corpus. It is possible that smaller groups exist with subtle language characteristics, which may be exploited with separate strategies. As a summary, we can conclude that understanding the writer motivation of any quotation, and in the broad area of opinion mining, is a key factor for the success of automatic classification.

References

- M. Al Masum Shaikh, H. Prendinger, and M. Ishizuka. 2007. Emotion sensitive news agent: An approach towards user centric emotion sensing from the news. In *Proceedings of the IEEE/WIC/ACM International Conference on Web Intelligence (WIC 2007)*, pages 614–620.
- A. Balahur and R. Steinberger. 2009. Rethinking sentiment analysis in the news: from theory to practice and back. In *Proceedings of the 1st Workshop on Opinion Mining and Sentiment Analysis*. University of Sevilla.
- A. Balahur, R. Steinberger, M. Kabadjov, V. Zavarella, E. van der Goot, M. Halkia, B. Pouliquen, and J. Belyaeva. 2010. Sentiment analysis in the news. In *Proceedings of the 7th Conference on International Language Resources and Evaluation*.
- Bomfim and Antulio N. 2000. Pre-announcement effects, news, and volatility: Monetary policy and the stock market. Technical report, Board of Governors of the Federal Reserve System.
- P. Chesley, B. Vincent, L. Xu, and R. Srihari. 2006. Using verbs and adjectives to automatically classify blog sentiment. In *Proceedings of the AAAI Symposium on Computational Approaches to Analyzing Weblogs (AAAI/CAAW 2006)*, pages 27–29.
- B. Drury, L. Torgo, and J.J Almedia. 2011. Classifying news stories to estimate the direction of a stock market index. In *Proceedings of the 3rd Workshop on Intelligent Systems and Applications*. <http://goo.gl/J7yv5>.
- V. Hatzivassiloglou and K.R. McKeown. 1997. Predicting the semantic orientation of adjectives. In *Proceedings of the 8th Conference on European Chapter of the Association for Computational Linguistics (EACL 1997)*, pages 174–181.
- V. Hatzivassiloglou and J. Wiebe. 2000. Effects of adjective orientation and gradability on sentence subjectivity. In *Proceedings of the 18th International Conference on Computational Linguistics (COLING 2000)*, pages 299–305.
- D.F. Larcker and A. Zakolyukina. 2010. Detecting deceptive discussions in conference calls.
- V. Lavrenko, M. Schmill, D. Lawrie, P. Ogilvie, D. Jensen, and J. Allan. 2000. Language models for financial news recommendation. In *Proceedings of the 9th International Conference on Information and Knowledge Management (CIKM 2000)*, pages 389–396.
- J. McManus. 1988. An economic theory of news selection. In *Annual Meeting for Education in Journalism and Mass Communication*.
- I. Mierswa, M. Wurst, R. Klinkenberg, M. Scholz, and T. Euler. 2006. Yale: Rapid prototyping for complex data mining tasks. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2006)*, pages 935–940.
- G. A. Miller. 1990. Wordnet: an on-line lexical database. *International Journal of Lexicography*, 3(4).
- B. Pang, L. Lee, and S. Vaithyanathan. 2002. Thumbs up?: Sentiment classification using machine learning techniques. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP 2002)*, pages 79–86.
- B. Pfahringer, H. Bensusan, and C. Giraud-Carrier. 2000. Meta-learning by landmarking various learning algorithms. In *Proceedings of the 17th International Conference on Machine Learning (ICML 2000)*, pages 743–750.
- Gerald Ratner. 2007. *The Rise and Fall... and Rise Again*. Wiley, J.
- C. Strapparava and R. Mihalcea. 2008. Learning to identify emotions in text. In *Proceedings of the 2008 ACM Symposium on Applied Computing (SAC 2008)*, pages 1556–1560.
- Tamura and Hiromichi. 2003. Individual-analyst characteristics and forecast error. *Financial Analysts Journal*.
- P.D. Turney. 2002. Thumbs up or thumbs down?: Semantic orientation applied to unsupervised classification of reviews. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics (ACL 2002)*, pages 417–424.
- J. Wiebe, T. Wilson, R. Bruce, M. Bell, and M. Martin. 2004. Learning subjective language. *Computational Linguistics*, 30(3):277–308.