

Linking News Sentiment to Microblogs: A Distributional Semantics Approach to Enhance Microblog Sentiment Classification

Tobias Daudert and Paul Buitelaar

Insight Centre for Data Analytics, Data Science Institute, National University of Ireland, Galway
firstname.lastname@insight-centre.org

Abstract

Social media's popularity in society and research is gaining momentum and simultaneously increasing the importance of short textual content such as microblogs. Microblogs are affected by many factors including the news media, therefore, we exploit sentiments conveyed from news to detect and classify sentiment in microblogs. Given that texts can deal with the same entity but might not be vastly related when it comes to sentiment, it becomes necessary to introduce further measures ensuring the relatedness of texts while leveraging the contained sentiments. This paper describes ongoing research introducing distributional semantics to improve the exploitation of news-contained sentiment to enhance microblog sentiment classification.

1 Introduction

In our increasingly digital society, we are subject to a deluge of unfiltered information not always objective or unbiased. The popularity of social media has made it a gateway to digital news content with 23% of the population in 2017 preferring this medium as a source of news¹. A particular case is Twitter and with the rise in popularity of this medium, short texts rich in information and/or sentiment are becoming a relevant source of information for the sharing of news stories (Mitchell and Page, 2015). However, traditional news are still important and at least as influential as digital media; in 2017, 32% of the people worldwide accessed digital news directly on a news website¹. In Twitter, over 85% of the retweets contain news mentions (Kwak et al., 2010). The diffusion of information is also crucial; people view what friends share leading to a fast diffusion of information with 75% of the total retweets occurring within a

¹<https://www.statista.com/chart/10262/selected-gateways-to-digital-news-content/>

day (Lerman and Ghosh, 2010; Kwak et al., 2010). This effect, combined with a higher perceived trust of shared information by friends, can lead to the construction of opinions based on already opinionated content (Zhao et al., 2011; Turcotte et al., 2015).

The importance of microblogs and news articles, their similar instantaneous availability, and their topic intersections lead us to explore how news articles and microblogs affect each other and, in detail, how the sentiments contained in both affect each other. This paper presents ongoing research dealing with this question and utilises distributional semantics, in detail, word embeddings, the cosine similarity, and the word mover's distance, to improve the modeling of the conveyance of news-contained sentiment on microblogs, aiming to enhance microblog sentiment classification.

2 Background

In the financial domain, prior research has shown the connection between sentiments and the market dynamics, exposing the financial domain as a relevant area for sentiment analysis in text (Van De Kauter et al., 2015; Kearney and Liu, 2014). Sentiments are contained in various forms of text including news and microblogs. It has been shown that positive news tend to lift markets whereas bad news tend to lower the markets (Schuster, 2003; Van De Kauter et al., 2015). Past research mainly focuses on news, particularly news titles (i.e headlines) (Nassiroussi et al., 2014; Kearney and Liu, 2014). However, not only sentiment contained in news is an important factor for the markets. For example, Bollen et al. (2011) linked changes in public mood to value shifts in the Dow Jones Industrial Index three to four days later. With an increasing magnitude of instantly available informa-

tion, factors affecting people’s sentiment rise. This includes other people’s textually-expressed sentiment since information is not always presented in a neutral manner. However, the relation between sentiments across different data sources, how they affect each other, and how this can be leveraged for sentiment classification has not been investigated yet.

2.1 Linking Sentiments Across Data Sources

Daudert et al. (2018) goes a step in this direction and exploits news sentiment to improve microblog sentiment classification. Their work utilises an **entity-based approach** which, given data annotated with sentiment, an entity e , and a period p , calculates the average sentiment for entity e in period p . The authors used a news dataset and calculated an average sentiment per company for news published between March 11th and 18th 2016 which was then used as additional information. Their assumption that within a certain period sentiments regarding the same entity should be similar across different data sources was examined. Using the average news sentiment performs well in periods when there is an overall sentiment other than neutral; in periods when the overall sentiment is neutral or balanced, a more sophisticated approach is needed. A neutral overall sentiment is achieved when positive and negative sentiment counteract with each other, independently of the number of news where each sentiment is expressed, whilst a balanced overall sentiment is achieved when the number of positive and negative news regarding a certain entity is similar. Given this, it becomes important to take a deeper look at news and microblogs as not all news are equally important to each microblog dealing with the same entity. Therefore, this research employs a distributional semantics approach to remove noise in terms of microblog-unrelated news sentiment although dealing with the same entity. To the best of our knowledge, only the previously mentioned work has started investigating the relations between the sentiments and leveraged them for microblog sentiment classification, hence, there is no research on the use of distributional semantics for sentiment linking. On the other hand, research targeting the field of semantic enrichment is available and it is particularly relevant when addressing the linking of news and microblogs (e.g. Guo et al. (2013); Wei et al. (2014); Abel et al. (2011);

Tsagkias et al. (2011)). Abel et al. (2011) suggests five different approaches of linking news to tweets: 1) a strict URL-based strategy, 2) a lenient URL-based strategy, 3) a bag-of-words strategy, 4) a hashtag-based strategy, and 5) an entity-based strategy. Strategy 5) comes close to what has been explored by Daudert et al. (2018) whereby our approach is inspired by 3), employing it as an add-on to 5). Other related research considering the combination of semantic similarity and sentiment analysis are (Tang et al., 2016; Poria et al., 2016). Poria et al. (2016) developed a Latent Dirichlet Allocation algorithm considering the semantic similarity between word pairs, instead of only utilising a word frequency measure, thus, capable of capturing opinions and sentiments that are implicitly expressed in a text and, overall, contributing to improved clustering. Tang et al. (2016) focused on learning word embeddings defined not only by context but also by sentiment. Their approach is able to better capture nearest neighboring vectors not only through their semantic similarity but also favoring the same sentiment polarity. This novel idea of utilising word embeddings to better capture polarity in documents was initially brought up by Maas et al. (2011).

The work described in this paper aims to address the existing knowledge gap concerning the application of distributional semantics for sentiment linking and assigning.

3 Methodology

The work performed is divided into two parts: the preparation of the data, and its use in a Machine Learning (ML) prediction model. Throughout this paper, we implement the methodology described by Daudert et al. (2018), utilising the same datasets (section 3.1) and experimental setup (section 3.4). We extend their previous work by improving the method to link a news sentiment to a microblog as well as to assign a news sentiment to a microblog (section 3.2).

The aim of this research is to explore the relation of sentiments between news and microblogs, hence, the linking of both data types becomes necessary. To fulfill this task, we leverage a microblog and a complimentary news dataset covering the same period and entities. For each microblog in the dataset, we model the sentiment conveyance between the news sentiment and the microblogs sentiment by assigning one news sentiment ac-

ording to each of the different methods as described in 3.2; these are then used as additional features for the Support Vector Machine (SVM).

This SVM is trained and tested with the datasets mentioned in section 3.1, aiming to explore whether the consideration of textual similarities for modeling the conveyed news sentiment can add value to the microblog sentiment classification. To investigate this, we compare a classification (1) purely based on microblog messages (table 2, *MT*) with (2) a classification based on microblog messages and entity-based news sentiment (table 2, *ES Agg.*), and (3) classifications based on microblog messages and context-based news sentiment (table 2, columns highlighted in gray).




Figure 1: Representation of dataset M and MRN, subset A and B. The two links represent matching entities in the datasets.

3.1 Data

This research makes use of two datasets: a microblog dataset (M) and a microblogs-related news dataset (MRN), represented in Figure 1. Dataset M contains microblogs from Twitter² as well as StockTwits³ and was initially created for the SemEval 2017 Task 5 - subtask 1 (Cortis et al., 2017); dataset MRN contains the news titles, urls, time and date, a sentiment score within the five classes [-1.0, -0.5, 0.0, 0.5, 1.0], and, if available, a description. All news in MRN are related to at least one microblog in dataset M. In total, MRN contains 106 news covering 18 unique entities in 463 microblogs (defined as subset A below). For dataset M, the sentiment scores are processed to cluster data in three classes by transforming sentiment scores above 0.0 to 1.0, and scores lower than 0.0 to -1.0. Moreover, two subsets of dataset M were created according to the microblogs’ re-

²<https://twitter.com>

³<https://stocktwits.com>

Type	Dataset M	Subset A	Subset B
Training	1,694	298	185
Test	794	165	92
Total	2,488	463	277

Table 1: Number of microblogs in dataset M, subset A, and subset B. Subset A and B are extracted from Dataset M.

lation to dataset MRN (see Table 1 and Figure 1). Subset A contains microblogs which have a relation to one or multiple news; subset B contains microblogs from subset A which are retrieved from Twitter. Subset B is necessary as dataset M contains StockTwits not specifically collected in the same period as the tweets. Figure 2 contains additional information regarding the annotation of both dataset as well as subsets.

3.2 Assigning a News Sentiment to Microblogs

All news in dataset MRN correspond to companies referred to in a minimum of one microblog in dataset M. With this information, our goal is to determine how to model the sentiment conveyance between the news-contained sentiment and each microblog given that news and microblogs might contain the same entities but not be vastly related. Considering the following example of two news articles, one about *Apple* and *Tim Cook’s* private life, and another one about *Apple* and the new *iPhone*, the latter one’s sentiment should have a higher impact on a microblog’s sentiment about Apple’s new products since they are more related. Using a purely entity based approach, both news articles would be linked to the microblog and the influence of both news on the assigned sentiment would be equal as they deal with *Apple*.

This work considers the assumption that “*within a certain period, sentiments regarding the same entity should be similar across different data sources*” (Daudert et al., 2018) and refines it with the assumption that *sentiments are particularly similar if the textual context is similar*. To lay the foundation for future research applications and to ensure a coherent understanding of the terminology applied throughout this work, we define core concepts as follows:

Linking - The linking of sentiment describes the creation of relations between sentiments, particularly their literal representations, by



Figure 2: Distribution of the annotated sentiment for dataset MRN, dataset M, subset A, and subset B. The sentiment is represented by s .

matching pieces of text according to predefined criteria such as entities, text intersections, or a degree of textual similarity. Hereby, we assume that linked sentiments are either influenced by the same cause or affecting each other.

Conveyance - The conveyance of sentiment describes the influence of the sentiment of one text on the sentiment of another. Sentiment is (indirectly) fully or partially transferred from a piece of text A to a piece of text B.

Assigning - The assigning of sentiment models the conveyance of sentiment from a text to another. Given two linked sentiments and the hypothesis that one is affecting the other, or both are affected by the same cause, we model the influence of text A's sentiment on text B's sentiment; improvements of this assignment can be measured by an enhanced sentiment detection for text B.

The aim behind this is the removal of noise in terms of microblog-unrelated news, although dealing with the same entity, as well as the reduction of the impact of less-related news on the assigned sentiment. To explore this, we compare four context-based approaches with the entity-based approach. The two context-based approaches employing a threshold for determining the relevance of a news to a microblog's sentiment (approach 1 and 3) aim at improving the sentiment linking since they fully discard news below a certain similarity value. The remaining two context-based approaches using a weighting scheme are reducing the impact of less relevant news on a microblog's sentiment and are, hence, aiming at improving the assigning of sentiment. This occurs in multiple steps: First, URLs in microblogs as well as news titles and descriptions are removed. Second, microblogs are tokenised employing the NLTK TweetTokenizer (Bird and Loper, 2004); news titles and descriptions are tokenized using the Stanford CoreNLP Tokenizer (Manning et al., 2014).

	Features	MT	ES Agg.	TS Thr.	TS Wgt. Agg.	WMD-S Thr. Agg.	WMD-S Wgt. Agg.
	Measure						
Dataset M	Micro F1-Score	0.8048	0.8060	0.8073	0.8060	0.8060	0.8060
	Macro F1-Score	0.6349	0.6357	0.6369	0.6357	0.6357	0.6357
	Weighted F1-Score	0.8018	0.8030	0.8044	0.8030	0.8030	0.8030
	Euclidean Distance	23.9165	23.8328	23.7487	23.8328	23.8328	23.8328
	Mean Error Squared	0.7204	0.7154	0.7103	0.7154	0.7154	0.7154
Subset A	Micro F1-Score	0.6485	0.6545	0.6606	0.6545	0.6485	0.6485
	Macro F1-Score	0.5547	0.5583	0.5633	0.5597	0.5547	0.5547
	Weighted F1-Score	0.64167	0.6471	0.6539	0.6485	0.6416	0.6416
	Euclidean Distance	14.9332	14.7986	14.6629	14.7986	14.9332	14.9332
	Mean Error Squared	1.3515	1.3273	1.3030	1.3272	1.3515	1.3515
Subset B	Micro F1-Score	0.7283	0.7283	0.7283	0.7391	0.75	0.7391
	Macro F1-Score	0.6917	0.6917	0.6917	0.701	0.7077	0.701
	Weighted F1-Score	0.7241	0.7241	0.7241*	0.7363	0.7463	0.7363
	Euclidean Distance	9.8489	9.8489	9.8489	9.6437	9.434	9.6437
	Mean Error Squared	1.0543	1.0543	1.0543	1.0109	0.9674	1.0109

Table 2: Scores obtained by the SVM model for dataset M, subset A, and B. MT abbreviates the message text, ES the entity-based news sentiment, TS the news title and description text similarity, WMD-S the word mover’s distance similarity. Thr. represents threshold, Wgt. weighted, and Agg. aggregated. A p-value < 0.01 is achieved for all models with the exception of *TS Thr.* on subset B (marked with *) which achieves a p-value < 0.05. The classifications based on microblog messages and context-based news sentiment are represented in gray (columns 4-7).

We choose different tokenizers for microblogs and news as the TweetTokenizer is specifically made for microblogs while news require a tokenizer adapted to a different structure and length. Third, we convert the Stanford GloVe Twitter model (Pennington et al., 2014) to Word2Vec (Mikolov et al., 2013a) and obtain the word embeddings. Having the word embeddings for microblogs and news in place, the subsequent processing varied depending on the context-based approach.

3.3 Context-based Approaches

We define **context-based** as an approach which utilises the textual similarity between two data artifacts as a factor to modify the sentiment of one of these, aiming at the generation of a sentiment to be assigned for the other artifact, necessary to model the sentiment conveyance.

In this work, we use microblog messages and a concatenation of the news titles and descriptions, if available, as our textual information. We then measure the textual similarity and utilise it as a factor to modify the news sentiment and subsequently generate the news sentiment to be assigned (NSTBA). This generated sentiment is then applied to model the sentiment conveyance be-

tween a news and a microblog.

$$NSTBA_m = \frac{s(n_1)+s(n_2)}{2} \quad (1)$$

The first context-based approach generates the NSTBA as an average of the sentiments of the microblog-related news articles. Document embeddings are retrieved for each microblog and news by averaging the word embeddings (Kartsaklis, 2014). We employ the cosine similarity as measure since vector offsets have been shown to be effective (Mikolov et al., 2013b). To be considered as context-related, a cosine similarity of at least 0.5 is required. For example, if two news articles (n_1, n_2) are context-related to microblog m , the two news sentiments (s) are added together and then divided by 2.

$$NSTBA_m = \frac{s(n_1)*sim(n_1,m)+s(n_2)*sim(n_2,m)}{2} \quad (2)$$

In contrast, the second context-based approach does not exclude relations with a cosine similarity lower than 0.5 but it uses the similarity score as a weighting factor multiplying it with the respective news sentiment score. Thus, an average of the similarity-weighted sentiments of the related-news is created. As an example, if two news articles (n_1, n_2) are context-related to microblog m ,

each news sentiment $s(n_x)$ is multiplied with the respective similarity (sim) score of n_x and m and then divided by 2. The NSTBA is then aggregated into the classes [-1.0, 0.0, 1.0] as this enhanced the results.

The third approach utilises the word mover’s distance (WMD) as described in (Kusner et al., 2015). We choose the WMD as it is a promising, recently developed function to measure the dissimilarity between two text documents. In our data, the WMDs d are within the range of [3.5, 9.5]. In spirit of equation 1, we use a threshold of 6.5 which is located halfway between both turning points as a requirement to be considered as context-related. As previously, the NSTBA has been aggregated into three classes.

The fourth approach is also based on the WMD. Since the WMD is not a similarity score but a distance theoretically ranging from 0 to unlimited, we transformed it into a similarity score (WMD-S). For WMDs ranging between [3.5, 9.5] in our data, we converted them into a similarity score within [0, 0.955] using the following formula:

$$sim(d) = 1 - \frac{1}{-0.1(d-9.5)^3+1} \quad (3)$$

Initially, we also experimented with other functions such as $1 - d/9.5$, however, function 3 represented a better approximation of a similarity score for our data. First, word embeddings are used to create the WMD between each microblog and news. Then, this distance is transformed into a similarity score using the formula above. Third and in the spirit of equation 2, news sentiments are weighted with the WMD-based similarity score. However, here we also aggregated the NSTBA.

3.4 Experimental Setup

For consistency, we utilise a similar setup to Daudert et al. (2018) for the preprocessing of the microblog texts, as well as for the SVM, and performance measures. The preprocessing steps are as follows:

1. URLs were replaced with $\langle url \rangle$
2. Numbers were replaced with $\langle number \rangle$
3. With *WORD* representing the original hashtag:
 - (a) hastags in upper case were replaced with $\langle hashtag \rangle WORD \langle allcaps \rangle$
 - (b) the remaining cases were replaced with $\langle hashtag \rangle WORD$

4. Smileys and emoticons were replaced with a description (e.g. ☺ becomes *slightly_smiling_face*)⁴

The processed text was then transformed into a unigram tf-idf representation.

The SVM model is trained and tested in six distinct approaches whereby approach three to six utilise different methods to model the context-based news sentiment: (1) a feature matrix representing the microblog messages; (2) a feature matrix representing the microblog messages enriched with the assigned entity-based news sentiment for each microblog, and (3)-(6) a feature matrix representing the microblog messages enriched with the assigned context-based news sentiment for each microblog. We chose to balance the class weight to get as close as possible to a *neutral sentiment* setting; the iterations are set to 500 and the random state to 42.

To test for statistical significance of the models, we apply a permutation test under the null hypothesis that the model has no effect in microblog sentiment classification (Ojala and Garriga, 2010).

4 Results

Table 2 shows the classification results on dataset M, subset A, and subset B. Although the use of an entity-based sentiment is already beneficial to the results, the addition of textual similarity measures further improves them. As the table shows, utilising context-based approaches to influence to-microblogs-assigned news sentiments enhances all measures in comparison to only using an entity-based average news sentiment. The weighted F1-Score for dataset M is increased by 0.17% and the Euclidean distance is decreased by 7.04%. In comparison to only using the message text (*MT*), the same scores are improved by 3.13% and 13.99%. For the subsets A and B the weighted F1-Score increases by 1.06% and 3.07%, and the Euclidean distance is decreased by 1.82% and 8.25%, respectively. For subset A, in contrast to only using *MT*, the weighted F1-Score and Euclidean distance are improved by 1.91% and 3.59%. This suggests the benefit of applying distributional semantics to the linking and assigning of news sentiment to microblogs, shown by the improvement on microblog sentiment classifica-

⁴<http://www.unicode.org/emoji/charts/full-emoji-list.html>

tion. Additionally, all scores improve on dataset M although only around 18.6% of the microblogs in the dataset are related to news. Surprisingly, utilising WMD-S improves all measures for subset B, whereas the cosine similarity between the document embeddings, together with the application of a threshold of 0.5, delivers the best results for dataset M and subset A. Furthermore, our approach outperforms the best score achieved in the SemEval 2017 Task 5 - Track 1 competition in which microblog sentiment analysis on a continuous scale was performed. Although our focus is to show the benefit of leveraging sentiment across news and microblogs, classifying the sentiment into 3 classes, our model reaches a cosine similarity of 0.869 on dataset M (table 2, column *TS Thr.*) whereas Jiang et al. (2017) reach a cosine similarity of 0.778.

5 Conclusion and Future Work

In this work, we utilise distributional semantics to model the conveyance of sentiment between news and microblogs. The achieved results suggest the benefit of using textual similarities and word embeddings to enhance the sentiment linking and assigning, culminating in an improved microblog sentiment classification. Our contributions are threefold: First, we present novel research utilising distributional semantics, specifically, word embeddings, the cosine similarity, and the word mover’s distance, for the linking and assigning of news-contained sentiment to microblogs; second, we explore the use of the word mover’s distance as similarity measure and; third, we suggest the benefit of leveraging news sentiment together with similarity methods for microblog sentiment classification. Comparing the additional use of an entity-based news sentiment with only the microblog text as features (columns *MT* versus *ES Agg.*), our results show an improvement on microblog sentiment classification on dataset M and subset A, while achieving a $p\text{-value} < 0.01$. In case of subset B, which has the most related news but the least news in quantity, the performance remains unchanged (columns *MT* versus *ES Agg.*). However, models utilising context-based news sentiment for an enhanced sentiment linking and assigning (columns *TS Thr.* and *WMD-S Thr. Agg.*) improve the performance for subset B and also reach the best scores for all three datasets. This suggests that applying distributional semantics is

particularly fruitful when entity-based news sentiments have less impact on the sentiment analysis on microblogs; this can be true in three cases:

1. The overall sentiments are neutral or balanced. We balanced all sentiment classes, however, the classifiers trained on context-based sentiment outperform the one trained on average entity-based news sentiment.
2. Only sparse related news exist. A classifier utilising the average entity-based sentiment as features achieves better results for dataset M and subset A than one with only the message text as features, however, on the smaller subset B this does not occur. Furthermore, when context-based sentiment is used as feature, the improvement on subset B becomes the largest. This suggests that each misleading news sentiment, present on dataset M and subset A, would have a noticeable impact on the results.
3. Related news are noisy and contain, apart from matching entities, unrelated information. Nonetheless, training our classifier on context-based sentiment outperforms the one trained on the average entity-based sentiment, suggesting that more-related news have a higher influence.

As future work, we aim to create a larger dataset, referring to a single defined period, linking microblogs and news. In addition, hybrid models taking into account not only a threshold for discarding noise but also a weighting scheme could potentially improve the classification. In this paper, we utilise the word mover’s distance and the cosine similarity to measure the similarity between two texts, however, other potentially adequate methods for this task still require exploration.

Acknowledgments

This publication has emanated from research conducted with the financial support of Science Foundation Ireland (SFI) under Grant Number SFI/12/RC/2289, co-funded by the European Regional Development Fund.

References

- Fabian Abel, Qi Gao, Houben Geert-Jan, and Tao KE. 2011. Semantic Enrichment of Twitter Posts for User Profile Construction on the Social Web. In *Proceedings of the Extended Semantic Web Conference*, pages 375 – 389. Springer, Berlin, Heidelberg.
- Steven Bird and Edward Loper. 2004. NLTK: The Natural Language Toolkit. In *Proceedings of the ACL 2004 on Interactive poster and demonstration sessions*, volume 1, pages 31 – 34, Morristown, NJ, USA. Association for Computational Linguistics.
- Johan Bollen, Huina Mao, and Xiaojun Zeng. 2011. Twitter mood predicts the stock market. *Journal of Computational Science*, 2(1):1–8.
- Keith Cortis, Andre Freitas, Tobias Daudert, Manuela Huerlimann, Manel Zarrouk, Siegfried Handschuh, and Brian Davis. 2017. SemEval-2017 Task 5: Fine-Grained Sentiment Analysis on Financial Microblogs and News. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 519–535.
- Tobias Daudert, Paul Buitelaar, and Sapna Negi. 2018. Leveraging News Sentiment to Improve Microblog Sentiment Classification in the Financial Domain. In *Proceedings of the First Workshop on Economics and Natural Language Processing*, pages 49–54, Melbourne, Australia. Association for Computational Linguistics.
- Weiwei Guo, Hao Li, Heng Ji, and Mona Diab. 2013. Linking Tweets to News: A Framework to Enrich Short Text Data in Social Media. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*, pages 239–249.
- Mengxiao Jiang, Man Lan, and Yuanbin Wu. 2017. ECNU at SemEval-2017 Task 5: An Ensemble of Regression Algorithms with Effective Features for Fine-Grained Sentiment Analysis in Financial Domain. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 888–893.
- Dimitri Kartsaklis. 2014. Compositional Operators in Distributional Semantics. *Springer Science Reviews*, 2(1-2):161–177.
- Colm Kearney and Sha Liu. 2014. Textual sentiment in finance: A survey of methods and models. *International Review of Financial Analysis*, 33:171–185.
- Matt J Kusner, Yu Sun, Nicholas I Kolkin, and Kilian Q Weinberger. 2015. From Word Embeddings To Document Distances. In *Proceedings of The 32nd International Conference on Machine Learning*, volume 37, page 957–966.
- Haewoon Kwak, Changhyun Lee, Hosung Park, and Sue Moon. 2010. What is Twitter, a Social Network or a News Media? In *Proceedings of the International World Wide Web Conference Committee (IW3C2)*, pages 591 – 600.
- Kristina Lerman and Rumi Ghosh. 2010. Information Contagion: an Empirical Study of the Spread of News on Digg and Twitter Social Networks. In *Proceedings of the Fourth International AAAI Conference on Weblogs and Social Media*, pages 90 – 97.
- Andrew L Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. Learning Word Vectors for Sentiment Analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 142–150.
- Christopher Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven Bethard, and David McClosky. 2014. The Stanford CoreNLP Natural Language Processing Toolkit. In *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 55 – 60.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. Efficient Estimation of Word Representations in Vector Space. *arXiv preprint arXiv:1301.3781*, pages 1–12.
- Tomas Mikolov, Wen-tau Yih, and Geoffrey Zweig. 2013b. Linguistic regularities in continuous space word representations. *Proceedings of NAACL-HLT*, (June):746–751.
- Amy Mitchell and Dana Page. 2015. The Evolving Role of News on Twitter and Facebook. Technical report, pewresearch.org.
- Arman Khadjeh Nassirtoussi, Saeed Aghabozorgi, Teh Ying Wah, and David Chek Ling Ngo. 2014. Text mining for market prediction: A systematic review. *Expert Systems with Applications*, 41(16):7653–7670.
- Markus Ojala and Gemma C. Garriga. 2010. Permutation Tests for Studying Classifier Performance. *Journal of Machine Learning Research*, 11(June):1833–1863.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. GloVe : Global Vectors for Word Representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, page 1532–1543.
- Soujanya Poria, Iti Chaturvedi, Erik Cambria, and Federica Bisio. 2016. Sentic LDA: Improving on LDA with semantic similarity for aspect-based sentiment analysis. In *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*, volume 201, pages 4465–4473. IEEE.
- Thomas Schuster. 2003. Meta-Communication and Market Dynamics. Reflexive Interactions of Financial Markets and the Mass Media. *SSRN eLibrary*, (July).

- Duyu Tang, Furu Wei, Bing Qin, Nan Yang, Ting Liu, and Ming Zhou. 2016. Sentiment Embeddings with Applications to Sentiment Analysis. *IEEE Transactions on Knowledge and Data Engineering*, 28(2):496–509.
- Manos Tsagkias, Maarten de Rijke, and Wouter Weerkamp. 2011. Linking online news and social media. In *Proceedings of the fourth ACM international conference on Web search and data mining - WSDM '11*, pages 565–574.
- Jason Turcotte, Chance York, Jacob Irving, Rosanne M. Scholl, and Raymond J. Pingree. 2015. News Recommendations from Social Media Opinion Leaders: Effects on Media Trust and Information Seeking. *Journal of Computer-Mediated Communication*, 20(5):520 – 535.
- Marjan Van De Kauter, Diane Breesch, and Veronique Hoste. 2015. Fine-grained analysis of explicit and implicit sentiment in financial news articles. *Expert Systems with Applications*, 42(11):4999–5010.
- Zhongyu Wei, Hong Kong, and Wei Gao. 2014. Utilizing Microblogs for Automatic News Highlights Extraction. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics*, pages 872 – 883.
- Wayne Xin Zhao, Jing Jiang, Jianshu Weng, Jing He, Ee-Peng Lim, Hongfei Yan, and Xiaoming Li. 2011. Comparing Twitter and Traditional Media Using Topic Models. In *Proceedings of the European Conference on Information Retrieval*, pages 338–349. Springer, Berlin, Heidelberg.