

Unsupervised Topic-Specific Domain Dependency Graphs for Aspect Identification in Sentiment Analysis

Sarah Kohail

Technische Universität Darmstadt,

Language Technology Group

Department of Computer Science

kohail@lt.informatik.tu-darmstadt.de

Abstract

We propose to model a collection of documents by means of topic-specific domain dependency graphs (DDGs). We use LDA topic modeling to detect topics underlying a mixed-domain dataset and select topically pure documents from the collection. We aggregate counts of words and their dependency relations per topic, weigh them with Tf-Idf and produce a DDG by selecting the highest-ranked words and their dependency relations. We demonstrate an implementation of the approach on the task of identifying product aspects for aspect-oriented sentiment analysis. A large corpus of Amazon reviews is used to identify product aspects by applying syntactic filtering to the DDG. Evaluation on a small set of camera reviews demonstrate a good precision of our method. To our knowledge, this is the first method that finds product-class specific aspects in multi-domain collections in an unsupervised fashion.

1 Introduction

Cohesion is reflected by grammatical and semantic relationships between lexical items, and links sentences together to form texts (Halliday and Hasan, 1976). These relationships contribute to the overall meaning of the text and maintain the inter-sentence and intra-sentence cohesive structure. Representations, such as graph-based have shown a potential ability to hold and understand these relationships, and facilitate knowledge extraction by enabling a variety of analysis processes (Radev and Mihalcea, 2008).

Recently, a large body of work has been devoted to applying graph or network-based methods to Natural Language Processing (NLP) problems, including, but not limited to, dependency

parsing (Tzouridis and Brefeld, 2013) to semantic annotation (Nivre and McDonald, 2008) to text summarization (Vidal et al., 2014) and information retrieval (Blanco and Lioma, 2012). In this paper, we present a generic graph-based method and apply it to identify product aspects for sentiment analysis.

E-commerce and social media technologies have become an excellent platform for a huge number of users to share and explain their opinions online. Websites (e.g., amazon.com, flipkart.com), allow users to post and read reviews about various services and products. Such reviews are important for customers to make a purchase choice, as well as for organizations to monitor and improve their products and reputation. However, user-generated reviews are unstructured and noisy. In the past few years, there has been a significant body of work that adopts NLP tools to better understand, analyze and process arguments and opinions from various types of information in user-generated reviews. Such efforts have come to be known as sentiment analysis or opinion mining, see (Liu, 2012) for a survey.

Sentiment analysis and opinion mining have been investigated on the document level, the sentence level and the aspect level (Liu, 2012). Aspect-level sentiment analysis performs fine-grained analysis by extracting or identifying the aspects of entities and the sentiment expressed toward each extracted aspect. For example, a review of a camera is likely to discuss distinct aspects like zoom, lens, resolution, battery life, price, and memory. In exploring the problem of aspect-based sentiment analysis, we distinguish between two terms "aspect identification" and "aspect extraction". Aspect extraction focuses on finding the aspects offsets in a given text reviews, while identification define the list of aspects of a certain entity.

The aim of this paper is to propose an unsupervised generic method to model a multi-domain

document collection by the means of domain dependency graphs (DDGs). An implementation of our method is applied to solve the aspect identification task from a large set of Amazon product reviews. The obtained graphs are used to improve the overall understanding of opinion patterns and to distinguish the most effective aspects for different product categories. Our method is completely unsupervised and needs no labeled training data or previous knowledge about the domains, and follows the Structure Discovery paradigm (Biemann, 2011). The remainder of this paper is organized as follows: Section 2 discusses related works. Section 3 describes the proposed solution. Section 4 presents and discusses our experimentation results and evaluation, followed by conclusions and future work in the last section.

2 Related Work

Graph theory has been widely used by many approaches in the field of natural language processing, text visualization and open information extraction (Koopman et al., 2012; Tzouridis and Brefeld, 2013), see (Mihalcea and Radev, 2011) for a survey. The most closely related work to our approach is (Stanovsky et al., 2014). It outlines Proposition Knowledge Graphs for information discovery. The utility of these knowledge graphs for structured queries, summarization and faceted search have been demonstrated.

In the field of sentiment analysis, graph-based approaches have been introduced to detect subjectivity (Esuli and Sebastiani, 2007; Wiebe and Mihalcea, 2006; Yu et al., 2011) or measure sentiment similarity between reviews (Goldberg and Zhu, 2006). Several methods were proposed to identify product aspects from reviews by selecting highly frequent nouns as product features (Blair-Goldensohn et al., 2008; Hu and Liu, 2004). For each detected noun, the sentiment regarding this noun is judged by its nearest adjacent adjective opinion word. However, the limitation of these methods is that many frequent noun phrases that may not represent product aspects are retrieved.

Recent research concentrates more on defining opinion patterns and relating aspects with their appropriate opinion words. Methodologies proposed in this area learn rules and templates from fully labeled data, and then use them later to detect aspects in an unlabeled dataset (Jin et al., 2009; Yu et al., 2011). Semi-supervised approaches try to re-

duce the amount of manual labeling by expanding a small seed set of labeled examples. Although these methods have been applied successfully in specific domains, sentiment classification is sensitive to the domain of the training data and extensive annotation for a large set of data for every single domain has to be carried out, which is not practically feasible (Vázquez and Bel, 2013).

Efforts for cross-domain sentiment analysis apply domain adaptation by limiting the set of features to those that are domain independent (Jakob and Gurevych, 2010; Li et al., 2012; Remus, 2012). An issue with these methods is that words and phrases used for expressing opinions can differ considerably from one domain to another.

3 Methodology

The purpose of this work is to advance understanding of a specific domain from mixed-domain documents by building compact directed DDGs. DDG aggregates individual dependency relations between domain-specific content words for a single topic. It gives a good visualization and summarization to a certain domain, and facilitate information and relation extraction. In this paper, we demonstrate the usage of DDGs for product aspects identification.

We summarize the methodology as follows: after preprocessing the text, we applied LDA topic modeling to discover underlying topics in a collection of textual data, and calculate a probabilistic topic distribution to select the most related phrases to each topic. POS tagging and dependency parsing were used then to select essential domain-specific phrases and content words. Finally, we build aggregate DDG per topic from the dependency parses, and use Tf-Idf and word frequency measures to weight the graph nodes and edges. A detailed discussion of our approach is given in the next section.

3.1 Dataset Preprocessing and Topic Modeling

Preprocessing includes filtering stop words, very short documents and documents with low frequency words. We perform word tokenization, and Latent Dirichlet Allocation (LDA) is then applied to extract dominant topics behind corpus of documents (Blei et al., 2003). LDA is a probabilistic graphical model that treats document as a multinomial distribution of topics, and each topic

is a multinomial distribution of words. LDA is completely unsupervised and requires no human annotation, but the user has to provide the number of topics n . We use the implementation provided by (Phan and Nguyen, 2007). We perceive all texts belonging to one topic i as one document d_i , where $i \in \{0, \dots, n\}$. The terms "domain" and "topic" are used interchangeably throughout the text.

3.2 Segmentation and Preprocessing

We use the vocabulary distribution of the documents produced by LDA to find a collection of topically pure documents. We retain only documents that have a single dominating topic, which covers at least 60% of the document¹. This step is significant to eliminate documents that contain too much noise or are too general to be characterize a specific topic. We then perform sentence segmentation² followed by POS tagging and collapsed dependency parsing³ (de Marneffe et al., 2006). The output from this step is important for generating syntactic features which will be used later to filter DDGs and extract topically pure relations.

3.3 Filtering Non-Content Words

For each document d_i , collapsed dependency document is generated. It includes a set of directed typed dependency relations R_{ijk} between a head word w_{ij} and a modifier word w_{ik} . As non-content words do not contribute as much information about a specific topic, we only retain relations between content words, i.e. (common and proper) nouns, adjectives, verbs and adverbs. From this step, the work followed is done completely on collapsed dependency documents.

3.4 Term Frequency-Inverse Document Frequency (Tf-Idf)

Tf-Idf is a standard term weighting method based on their importance within a document. The core idea behind Tf-Idf is: a word j w_{ij} in document i is more relevant as a keyword for d_i if it appears many times in d_i and very few times or none in other set of documents in a corpus D . Tf-Idf is ex-

pressed by the following equation:

$$Tf-Idf(w_{ij}, d_i, D) = Tf(w_{ij}, d_i) \times Idf(w_{ij}, D) \quad (1)$$

where Tf is the number of times that word w occurs in document d and Idf is calculated by dividing the total number of documents in a corpus, which is the number of topics n , by the number of documents containing the word w in a set of documents D .

Tf-Idf is calculated in three levels of granularity:

1. Word level: for each word w_{ij} in d_i , we calculated Tf-Idf using Equation 1.
2. Pair level: for each pair of words w_{ij} and w_{ik} in d_i , occurred together in a typed dependency relation R_{ijk} , we calculated Tf-Idf using the following equation:

$$Tf-Idf(w_{ij}w_{ik}, d_i, D) = Tf(w_{ij}w_{ik}, d_i) \times Idf(w_{ij}w_{ik}, D) \quad (2)$$

w_{ij} and w_{ik} represents the j^{th} and k^{th} words in document i . Order of words w_{ij} and w_{ik} within the relation is not considered at this level.

3. Relation level: for each typed dependency relation R_{ijk} in d_i between two words w_{ij} and w_{ik} , we calculate Tf-Idf using the following equation:

$$Tf-Idf(R_{ijk}w_{ij}w_{ik}, d_i, D) = Tf(R_{ijk}w_{ij}w_{ik}, d_i) \times Idf(R_{ijk}w_{ij}w_{ik}, D) \quad (3)$$

3.5 Domain Dependency Graphs (DDGs)

DDGs are directed graph with labeled nodes and labeled edges. For each document d_i , DDG_i is constructed by aggregating individual dependency relations between domain-specific content words. $DDG_i = \{V_i, E_i\}$, where nodes represent words, that is $V_i = \{w_{ij} \mid w_{ij} \in d_i, Tf-Idf(w_{ij}, d_i, D) \geq \alpha_1, Tf(w_{ij}) \geq \alpha_2\}$, and edges E_i connect content words by the means of dependency relations. $E_i = \{(w_{ij}, w_{ik}) \mid w_{ij}, w_{ik} \in d_i, Tf-Idf(w_{ij} w_{ik}, d_i, D) \geq \beta_1, Tf(w_{ij} w_{ik}) \geq \beta_2, Tf-Idf(R_{ijk} w_{ij} w_{ik}, d_i, D) \geq \lambda_1, Tf(R_{ijk} w_{ij} w_{ik}) \geq \lambda_2\}$.

Thresholds, α_1 , α_2 , β_1 , β_2 , λ_1 , λ_2 are defined

¹Threshold was determined in preliminary experiments

²Using lt.seg script from <https://github.com/tudarmstadt-lt/lt.core/>

³We use the Stanford Natural Language Processing tools <http://nlp.stanford.edu/software/>

by the user, and edges are labeled by the frequency and the type of dependency relation between words. Using Tf-Idf for weighting words and relations, have proven a potential ability to highlight a large set of domain-specific words and relations as will be demonstrated in the next section.

3.6 Extracting Domain Dependency Words and Relations - Application

We apply our generic approach to identify opinion phrases, and aspects of products for the use in aspect-based sentiment analysis.

Figure 1 illustrates a snapshot from DDG for a topic that captures camera reviews. We use DDGs along with Tf-Idf weighting as an important input to distinguish most related domain specific words and relation patterns. We present below some words examples from the camera’s domain categorized by POS tags. All mentioned words are strongly related to camera domain and this proves the capability of Tf-Idf weighting in capturing potential domain specific words.

- **Adjectives:** digital, 50mm, focal, 200mm, optical, sharp, indoor, blurry, wide, prime, compact, chromatic.
- **Nouns:** lens, camera, canon, nikon, SLR, EF, shots, shutter, USM, telephoto, aperture, macro, flash, sigma, focus, pictures, zoom, tripod, powershot.
- **Verbs:** taking, focuses, capture, carry, photographing, fit, produce, cropping, adjust.

We highlight some opinion relations from Figure 1 in Table 1. The table shows dependency relation type R_{Camjk} , source word w_{Camj} , destination word w_{Camk} , relation frequency Tf and relation level $Tf-Idf$. We create DDGs for another 14 topics including: movies, coffee makers, electro-voice, shoes and footwear, hair products, food and baking machines, films, mp3 players, cars, TVs, mobiles, computers and perfumes. We observed that in all these graphs, opinions or relations between opinion word and opinion target, are mostly expressed with either adjectival modifier (amod) or nominal subject (nsubj). Thus, we will limit the identification of product aspects to these two dependency relations in our application.

On the basis of our analysis of DDGs and their parameters, and a list of about 6800 words positive and negative English opinion words⁴, we apply a

⁴English Opinion Lexicon <http://www.cs.uic.edu/~liub/FBS/sentiment-analysis.html#lexicon>

set of appropriate filters to DDG to extract opinion phrases. We filter out noun compounds relations, and words and relations below thresholds α_1 , α_2 , β_1 , β_2 , λ_1 , λ_2 . Either w_{ij} or w_{ik} should be in opinion words lexicon and relation which is either "amod" or "nsubj" is selected.

R_{Camjk}	w_{Camj}	w_{Camk}	Tf	$Tf-Idf$
amod	lens	fast	146	770.60
nsubj	great	lens	121	638.65
amod	picture	good	205	467.88
amod	images	sharp	116	451.45
nsubj	sharp	images	93	388.69
amod	photos	great	105	269.85
amod	picture	clear	84	241.93
nsubj	good	quality	142	50.85

Table 1: Opinion dependency relations from the camera topic.

4 Experiments

To evaluate our approach, we use an unlabeled version of Amazon dataset⁵ that has been commonly used in opinion mining research (Kiritchenko et al., 2014; Tutubalina, 2015). The corpus consists of ~ 35 million reviews (~ 18.4 million unique reviews), about ~ 2.5 million products from 28 different categories, up to March 2013. Reviews include product and user information, ratings, and a plain text review (McAuley and Leskovec, 2013).

In this work, we only use the plain text. We filter redundant reviews, reviews with less than 3 words and noisy reviews which contain smiley codes only or punctuations only, as we consider these not relevant for aspect identification. The final number of reviews we use to train the LDA model is ~ 13.93 million reviews. As we mentioned in Section 3.2, we use the LDA model to select topically pure reviews. This step reduces the number of reviews to ~ 1 million.

We experimentally determined a reasonable number of topics n to be 200, which is in line with other works using LDA for information extraction e.g. (Chambers and Jurafsky, 2011). Of the 200 topics we induced with LDA, we observed a large number of product-specific topics, as well as some mixed topics and spurious topics (Mimno et al., 2011). For this study, we proceed with selecting the 15 topics we mentioned in Section 3.6. To

⁵SNAP: Web data: Amazon reviews <https://snap.stanford.edu/data/web-Amazon.html>

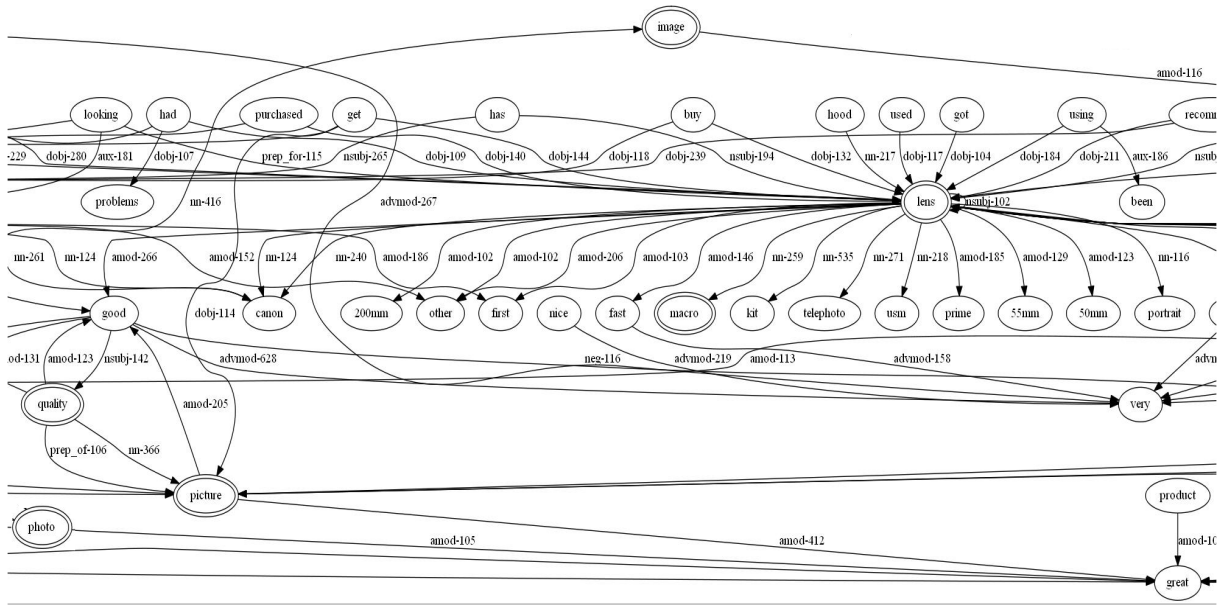


Figure 1: An excerpt from the automatically generated DDG of the camera review topic. Double lining for aspect nodes, and bold lines for connections between opinion words and aspects have been assigned manually. Only most frequent relations are shown for the purpose of presentation.

test the performance of our proposed approach, we compare our results to those obtained using DDG without Tf-Idf filtering, i.e. $\alpha_1 = \beta_1 = \lambda_1 = 0$. We evaluate the identification of aspects manually by human judgment: We order the identified relations from both Tf-Idf-based filtered DDGs (as explained in Section 3.6) and frequency-based (FB) filtered DDGs according to relation frequency. For the top 50 unique aspects, we judge whether it is an aspect of the product category or not.

Table 2 shows the experimental results for 5 different product topics. The experimental results show that Tf-Idf filtering outperforms FB filtering in terms of the number of identified aspects and it has not been worse in any case. FB ranking tend to identify general aspects such as: price, shipping, quality, value, service and company. Ranking DDGs by the means of Tf-Idf weights, gives our method the ability to detect detailed domain aspects, which is clearly evident in the cars topic in Fig. 2. The aspect identification method based on the DDG with Tf-Idf weighting identifies domain-specific aspects with an average accuracy of 53% across the five topics. When not using Tf-Idf weighting, the method achieves only an accuracy of 37%.

Our error analysis shows that most false positives by the Tf-Idf-based method consist of product domain-specific words that are not aspects.

Examples from cameras domain are: fast results, great job cheap camera, excellent choice, sharp razor, perfect bag, great portrait, advanced photographer, easy c330. On the other hand, frequency-based ranking provides general noisy errors like: problem only, buy great, complaint only, time hard, addition great, drawback only, light available, room enough.

To evaluate the identified aspects coverage for the aspects extraction task from a set of reviews, we manually annotated aspects in a set of 50 cameras reviews collected randomly from Amazon. Only explicit aspects are annotated. Implicit aspects are not annotated. In most of implicit aspect expressions, adjectives and adverbs are used to describe some specific attributes of entities, for example, expensive describes price, and heavy describes weight (Liu, 2012). We compared the annotated aspects against the 33 aspects for cameras domain listed in Table 2. Out of 183 annotated aspects in the 50 reviews, 115 aspects are extracted, approximately 63%, while 38 unique failed to be extracted. Most of missed aspects are contained in cameras reviews DDG before filtering. Changing the filtering parameters can help increasing the aspects coverage but may also increase the false positive rate.

In summary, our evaluation shows a clear improvement using Tf-Idf-based filtering over the

Category / Thresholds	Method	Ext. /50	Extracted Aspects	
			Common	Difference
Camera $\alpha_1: 100, \alpha_2: 180$ $\beta_1: 2, \beta_2: 2$ $\lambda_1: 7, \lambda_2: 5$	Tf-Idf-based	30	lens, pictures, shots, quality, images, photos, focus, light, depth, color, zoom, size, range, distortion, card, autofocus, speed.	tripod, resolution, controls, battery, mode, contrast, optics, flash, sharpness, software, screen, flexibility, distance.
	FB	20		price, value, capability.
TV $\alpha_1: 50, \alpha_2: 20$ $\beta_1: 1, \beta_2: 1$ $\lambda_1: 2, \lambda_2: 5$	Tf-Idf-based	22	cable, picture, quality, remote, setup, image.	system, audio, resolution, output, video, tuner, hdtv, quality, connection, capability, control, speakers, screen, model, component, connector.
	FB	13		price, sound, value, shipping, colors, monitor, pixels.
Computer $\alpha_1: 150, \alpha_2: 50$ $\beta_1: 2, \beta_2: 2$ $\lambda_1: 2, \lambda_2: 5$	Tf-Idf-based	29	card, software, memory, adapter, performance, setup, support, camera, driver, ram, disk, space, cable.	upgrade, programs, ports, system, processor, speed, motherboard, version, machine, units, USB, slots, OS, mouse, graphics, interface.
	FB	19		price, power, value, quality, shipping, case.
Mobile $\alpha_1: 50, \alpha_2: 20$ $\beta_1: 1, \beta_2: 1$ $\lambda_1: 5, \lambda_2: 1$	Tf-Idf-based	20	sound, keyboard, screen, price, reception, quality, size, case, camera, service, software.	pictures, apps, life, interface, looks, bluetooth, battery, version, calls.
	FB	18		card, program, version, design, charger, player, value.
Cars $\alpha_1: 20, \alpha_2: 5$ $\beta_1: 2, \beta_2: 1$ $\lambda_1: 5, \lambda_2: 1$	Tf-Idf-based	32	price, performance, exhaust, wiring, plugs, installation, power, length, kit, sound, shocks, sensors, ride, instructions, parts.	work, rumble, breaks, pads, muffler, replacement, wipers, harness, connectors, idle, engine, hitch, system, unit, lights, mileage, tensioner.
	FB	23		quality, shipping, value, struts, company, service, look, room.

Table 2: Manual evaluation for aspect identification on five different domains using DDG with Tf-Idf ranking and FB ranking. It shows the number of true identified aspects out of the top 50 frequent captured relations, common identified aspects along with the difference between the two methods. The first column shows the thresholds setting. For the frequency-based ranking method, $\alpha_1 = \beta_1 = \lambda_1 = 0$.

FB baseline. This, however, is only possible for mixed-domain document collections, as Idf for a single topic is not defined.

5 Conclusion

We have introduced a new generic approach to identify the most important concepts from multi-domain document collections. Using LDA, we provided a fully unsupervised framework for extracting dominant topics behind corpus of documents, while the DDG representation maintains the inter-topic cohesiveness. Tf-Idf ensures the extraction of highly domain-specific words and relations. We demonstrate the effectiveness of the proposed approach on the task of extracting product aspects for sentiment analysis. The comparison between the DDG method and a frequency-based ranking confirms the superiority of DDG in extracting domain-specific aspects. Evaluation of

DDG on a small set of cameras reviews resulted in a precision of $\sim 63\%$. This is the first approach, to our knowledge, for extracting product aspects from mixed-domain dataset, without the use of an external knowledge base or a training dataset.

In the future, we hope to advance our work by using DDGs to applying more advanced ranking and filtering techniques to DDGs such as centrality (Newman, 2010) or PageRank (Brin and Page, 1998) for node ranking. Collecting similarities to the existing list of aspects and grouping aspects using techniques from distributional semantics would improve the overall recall.

Acknowledgments

This research was supported by the DAAD. The authors sincere gratitude goes to Chris Biemann for his full supervision. Also, the author would like to thank the members of Language Technology group in TU Darmstadt and the anonymous reviewers who greatly refined the drafts of this paper.

References

- Chris Biemann. 2011. *Structure Discovery in Natural Language*. G. Hirst, E. Hovy and M. Johnson (Series Eds.): Theory and Applications of Natural Language Processing. Springer, Heidelberg Dordrecht London New York.
- Sasha Blair-Goldensohn, Kerry Hannan, Ryan McDonald, Tyler Neylon, George A Reis, and Jeff Reynar. 2008. Building a sentiment summarizer for local service reviews. In *WWW Workshop on NLP in the Information Explosion Era*, pages 14–23, Beijing, China.
- Roi Blanco and Christina Lioma. 2012. Graph-based term weighting for information retrieval. *Information retrieval*, 15(1):54–92.
- David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *the Journal of machine Learning research*, 3:993–1022.
- Sergey Brin and Lawrence Page. 1998. The anatomy of a large-scale hypertextual Web search engine. *Computer Networks and ISDN Systems*, 30(1-7):107–117.
- Nathanael Chambers and Dan Jurafsky. 2011. Template-based Information Extraction Without the Templates. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies - Volume 1*, HLT '11, pages 976–986, Portland, Oregon.
- Marie-Catherine de Marneffe, Bill MacCartney, and Christopher D. Manning. 2006. Generating typed dependency parses from phrase structure parses. In *Proceedings of LREC*, volume 6, pages 449–454, Genoa, Italy.
- Andrea Esuli and Fabrizio Sebastiani. 2007. Pageranking WordNet synsets: An application to opinion mining. In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics ACL-07*, volume 7, pages 442–431, Prague, Czech Republic.
- Andrew B Goldberg and Xiaojin Zhu. 2006. Seeing stars when there aren't many stars: graph-based semi-supervised learning for sentiment categorization. In *Proceedings of the First Workshop on Graph Based Methods for Natural Language Processing*, pages 45–52, New York, NY, USA.
- Michael AK Halliday and Ruqaiya Hasan. 1976. Cohesion in English. *Longman's, London*.
- Minqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 168–177, New York, NY, USA.
- Niklas Jakob and Iryna Gurevych. 2010. Extracting opinion targets in a single-and cross-domain setting with conditional random fields. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 1035–1045, Cambridge, Massachusetts.
- Wei Jin, Hung Hay Ho, and Rohini K Srihari. 2009. A Novel Lexicalized HMM-based Learning Framework for Web Opinion mining. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 465–472, Montreal, Quebec, Canada.
- Svetlana Kiritchenko, Xiaodan Zhu, Colin Cherry, and Saif Mohammad. 2014. NRC-Canada-2014: Detecting Aspects and Sentiment in Customer Reviews. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 437–442, Dublin, Ireland.
- Bevan Koopman, Guido Zuccon, Peter Bruza, Laurianne Sitbon, and Michael Lawley. 2012. Graph-based concept weighting for medical information retrieval. In *Proceedings of the Seventeenth Australasian Document Computing Symposium*, pages 80–87, Dunedin, New Zealand.
- Fangtao Li, Sinno Jialin Pan, Ou Jin, Qiang Yang, and Xiaoyan Zhu. 2012. Cross-domain co-extraction of sentiment and topic lexicons. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1*, pages 410–419, Jeju Island, Korea.
- Bing Liu. 2012. Sentiment analysis and opinion mining. *Synthesis Lectures on Human Language Technologies*, 5(1):1–167.
- Julian McAuley and Jure Leskovec. 2013. Hidden factors and hidden topics: understanding rating dimensions with review text. In *Proceedings of the 7th ACM conference on Recommender systems*, pages 165–172, Hong Kong, China.
- Rada Mihalcea and Dragomir Radev. 2011. *Graph-based natural language processing and information retrieval*. Cambridge University Press.
- David Mimno, Hanna Wallach, Edmund Talley, Miriam Leenders, and Andrew McCallum. 2011. Optimizing Semantic Coherence in Topic Models. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing, EMNLP '11*, pages 262–272, Edinburgh, United Kingdom.
- Mark Newman. 2010. *Networks: an introduction*. Oxford University Press.
- Joakim Nivre and Ryan Mcdonald. 2008. Integrating Graph-Based and Transition-Based Dependency Parsers. In *Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies (ACL-08: HLT)*, pages 950–958, Columbus, Ohio.
- Xuan-Hieu Phan and Cam-Tu Nguyen. 2007. GibbsLDA++: A C/C++ implementation of latent Dirichlet allocation (LDA). <http://gibbslda.sourceforge.net>.
- Dragomir R Radev and Rada Mihalcea. 2008. Networks and natural language processing. *AI magazine*, 29(3):16–28.
- Robert Remus. 2012. Domain adaptation using domain similarity-and domain complexity-based instance selection for cross-domain sentiment analysis. In *IEEE 12th International Conference on Data Mining Workshops (ICDMW 2012)*, pages 717–723, Brussels, Belgium.
- Gabriel Stanovsky, Omer Levy, and Ido Dagan. 2014. Proposition Knowledge Graphs. In *Proceedings of the First AHA!-Workshop on Information Discovery in Text*, pages 19–24, Dublin, Ireland.
- Elena Tutubalina. 2015. Target-Based Topic Model for Problem Phrase Extraction. In *Advances in Information Retrieval*, pages 271–277. Springer.

- Emmanouil Tzouridis and Ulf Brefeld. 2013. Learning the Shortest Path for Text Summarisation. In *The Fourth International Workshop on Mining Ubiquitous and Social Environments*, pages 45–57, Prague, Czech Republic.
- Silvia Vázquez and Núria Bel. 2013. A classification of adjectives for polarity lexicons enhancement. *arXiv preprint arXiv:1303.1931*.
- Juan C Vidal, Manuel Lama, Estefanía Otero-García, and Alberto Bugarín. 2014. Graph-based semantic annotation for enriching educational content with linked data. *Knowledge-Based Systems*, 55:29–42.
- Janyce Wiebe and Rada Mihalcea. 2006. Word sense and subjectivity. In *Proceedings of the 21st International Conference on Computational Linguistics and the 44th annual meeting of the Association for Computational Linguistics*, pages 1065–1072, Sydney, Australia.
- Jianxing Yu, Zheng-Jun Zha, Meng Wang, and Tat-Seng Chua. 2011. Aspect ranking: identifying important product aspects from online consumer reviews. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 1496–1505, Portland, Oregon.