FACHLITERATUR
EDITION ROMIOSINI
ΕΠΙΣΤΗΜΗ

ICGL12 | 12th INTERNATIONAL CONFERENC ON GREEK LINGUISTICS
16 – 19 SEPTEMBER 2015
FREIE UNIVERSITÄT BERLIN, CEMOG

Proceedings of the ICGL12

vol. 2

The International Conference on Greek Linguistics is a biennial meeting on the study and analysis of Greek (Ancient, Medieval and Modern), placing particular emphasis on the later stages of the language.

PROCEEDINGS OF THE ICGL12
ΠΡΑΚΤΙΚΑ ΤΟΥ ICGL12

Thanasis Georgakopoulos, Theodossia-Soula Pavlidou, Miltos Pechlivanos, Artemis Alexiadou, Jannis Androutsopoulos, Alexis Kalokairinos, Stavros Skopeteas, Katerina Stathi (Eds.)

PROCEEDINGS OF THE 12TH INTERNATIONAL CONFERENCE ON GREEK LINGUISTICS

ΠΡΑΚΤΙΚΑ ΤΟΥ 12ΟΥ ΣΥΝΕΔΡΙΟΥ ΕΛΛΗΝΙΚΗΣ ΓΛΩΣΣΟΛΟΓΙΑΣ

VOL. 2

Vertrieb und Gesamtherstellung: Epubli (www.epubli.de)
Satz und Layout: Rea Papamichail / Center für Digitale Systeme, Freie Universität Berlin
Gesetzt aus Minion Pro
Umschlaggestaltung: Thanasis Georgiou, Yorgos Konstantinou
Umschlagillustration: Yorgos Konstantinou

ISBN 978-3-946142-35-5
Printed in Germany

Online-Bibliothek der Edition Romiosini:
www.edition-romiosini.de

ΠΕΡΙΕΧΟΜΕΝΑ

OPINION MINING AND TARGET EXTRACTION IN GREEK REVIEW TEXTS

Maria Pontiki
Institute for Language and Speech Processing, Athena RIC, Greece
mpontiki@ilsp.gr

Περίληψη

Στο παρόν άρθρο παρουσιάζεται μια βασισμένη σε κανόνες γλωσσολογική προσέγγιση για την ανάλυση άποψης στην ελληνική γλώσσα. Δοθέντος ενός κειμένου που περιέχει κριτικές για κάποια συγκεκριμένη οντότητα (π.χ. ένα εστιατόριο ή ένα ξενοδοχείο), στόχος είναι ο εντοπισμός των απόψεων που εκφράζονται στο κείμενο, καθώς και των οντοτήτων-στόχων για τις οποίες εκφράζεται η κάθε άποψη. Για τον σκοπό αυτό, αναπτύχθηκε ένας αναλυτής άποψης που βασίζεται σε λεξικά και αποτελείται από μια σειρά αρθρωμάτων (modules) που περιλαμβάνουν κανόνες επιφανειακής συντακτικής δομής. Το βασικό πλεονέκτημα των γλωσσολογικών προσεγγίσεων είναι τα υψηλά ποσοστά ακρίβειας, τα οποία και επιβεβαιώνονται από τα αποτελέσματα της πειραματικής αξιολόγησης.

Keywords: sentiment analysis, opinion mining, opinion target extraction, customer reviews, rule-based approach, lexicon-based method, finite state transducers

1. Introduction

Mining opinions from user-generated content on the Web (e.g. review sites) has attracted significant commercial and academic interest given the variety of its practical applications and the research challenges it involves. Online reviews evolved in an important source of past consumption experiences (Yoo and Gretzel 2008) and, thus,

of opinionated user-generated content. Related research highlights the importance of online reviews for the consumer decision making process (e.g., when purchasing experiential goods like travel destinations, hotels, and restaurants (Liu 2015)), as well as for product sales (Duan et al. 2008)).

A fundamental problem in opinion mining is the extraction of the opinion targets (Pang and Lee 2008), namely identifying the entities, objects, or topics on which opinions are expressed. Opinion targets are important because without knowing what an opinion is about, the identified (positive or negative) sentiments within a sentence or a document are somewhat meaningless and of limited use. For example, restaurant customer reviews not only express the overall sentiment about a specific restaurant (e.g., (1)), but also opinions related to different aspects of the food (e.g., the quality and the price of the food in (2)), or the service (e.g. (3)).

(1) "Ωραίο μέρος."
"Nice place."
(2) "Τα κυρίως πιάτα ήταν πολύ νόστιμα αλλά και πολύ ακριβά επίσης!"
"The main courses were delicious but also very expensive!"
(3) "Ο σερβιτόρος ήταν πολύ αγενής."
"The waiter was very rude."

In this paper, we present a linguistically driven opinion mining method for the Greek language. The goal is, given a review text about a particular target entity (e.g., a restaurant), to identify and automatically extract all the opinion expressions and their targets along with the respective sentiment polarity label. For this purpose we have developed a rule-based opinion analyzer (i.e. a set of linguistic rules that model phrase-level opinion expressions). The analyzer relies on EvalLex (Pontiki et al. 2013), an Appraisal Theory (Martin and White 2005) grounded Lexicon for Evaluative Language that was manually compiled for the Greek language. In order to achieve better recall, we further expanded it with other types of sentiment vocabulary. Another significant contribution of this work is the creation of Greek datasets that contain human authored annotations of opinion expressions and targets.

The remainder of this paper is structured as follows. Section 2 provides an overview of the related research. The proposed method for opinion mining and target extraction is described in section 3. The evaluation results are presented and discussed in section 4. The paper concludes with some brief remarks on possible future directions.

2. Related Work

A variety of opinion mining methods that involve the extraction of opinion targets have been proposed (Hu and Liu 2004, Ding et al. 2008, Li et al. 2010, Popescu and Etzioni 2005). Opinion target approaches usually rely on opinion words (i.e. vocabulary used to express opinions) and exploit their relations and associations with the words they modify using syntactic-based approaches (Popescu and Etzioni 2005, Qiu et al. 2011), word alignment models (Liu et al. 2015) or hybrid (joint) methods (Liu et al. 2013). For example, Qiu et al. (2011) use a double-propagation method for the bidirectional transfer of sentiment values onto targets and back to unknown sentiment terms based on dependency relations. Opinion target extraction has also been addressed in the context of Aspect Based Sentiment Analysis (Liu and Zhang 2013, Pontiki et al. 2014, Pontiki et al. 2015).

Supervised approaches usually model target extraction as a sequential labelling task (Li et al. 2010, Jakob and Gurevych 2010, Ma and Wan 2010, Wu et al. 2009, Zhang et al. 2009) using Conditional Random Fields (Lafferty et al. 2001). Unsupervised approaches rely on a variety of rules and techniques. For example, Hu and Liu (2004) assume that the association between terms can predict features and opinion words and employ an association rule mining technique (Agrawal and Srikant 1994). Popescu and Etzioni (2005) extract product features and opinions from unstructured reviews based on syntactic patterns, while Carenini et al. (2005) employ a domain-dependent approach using a predefined taxonomy of product features. The few opinion mining methods that have been developed for the Greek language (Karampatsis 2014) or applied to it in the context of multilingual approaches (Solakidis et al. 2015), focus only on the overall sentiment classification of short text messages not involving the extraction of the respective opinion targets. In the work of Agathagelou et al. (2014) candidate opinion targets (e.g., nouns) are used along with an available list of opinion words in order to extract new opinion words based on a double propagation method.

In the work presented in this paper the goal is, given a review text about a particular target entity (e.g., a restaurant), to identify and automatically extract all the opinion expressions and all the opinion targets along with the respective sentiment polarity labels. We employ a rule-based approach similar to Qiu et al. (2011) with the following two differences. First, unlike the method of Qiu et al. (2011) that performs opinion lexicon expansion (i.e. identifies new opinion words) based on a seed opinion lexicon, our lexicon-based method detects expressed opinions by performing opinion detec-

tion and polarity disambiguation in the context around EvalLex's (Pontiki et al. 2013) entries. Second, the syntactic patterns in our method model shallow syntactic relations and do not make use of dependency relations.

3. Methodology

We have developed a precision-oriented opinion analyzer that, given a review text about a particular target entity (e.g., a restaurant), is designed to extract opinion expressions (i.e. instantiations of opinions expressed in a text) and their targets along with the respective polarity focusing only on positive and negative polarity. The method is lexicon-based and explores shallow syntactic relations between sentiment words (i.e. words that are used to express opinions, emotions, etc.) and sequences of Tokens-candidate opinion targets. In particular, the method relies on lexical resources (EvalLex and a manually built lexicon of Modifiers) and a four-phase grammar of linguistic rules (Opinion Grammar).

The opinion analyzer is a Finite State Transducers (FST) cascade implemented as a JAPE grammar (Cunningham et al. 2000) in the GATE framework. The input for the opinion analyzer is raw data. In a first phase the datasets are processed through a NLP pipeline that performs tokenization, sentence splitting, part-of-speech tagging, and lemmatization using the ILSP suite of NLP tools for Greek (Papageorgiou et al. 2002, Prokopidis et al. 2011). In a subsequent step, EvalLex is used to determine which tokens are included in the lexicon. If a token is recognized as a lexicon entry, then it was annotated with the EvalLex attributes. Then, sets of hand-crafted rules determine which spotted EvalLex's entries express sentiment towards specific targets based on lexico-syntactic patterns that model shallow syntactic relations. The overall workflow of our method is illustrated in Figure 1.

3.1. Datasets

The method was built based on the customer reviews datasets of (Pontiki et al. 2013). In particular, the "Restaurant Reviews" dataset consists of 1058 sentences from user generated reviews about specific restaurants (e.g., ethnic, fast food, gourmet, traditional Greek) manually collected from *ask4food.gr*. The "*Movie Reviews*" dataset contains 1007 sentences from user generated reviews about specific movies (e.g., comedies,

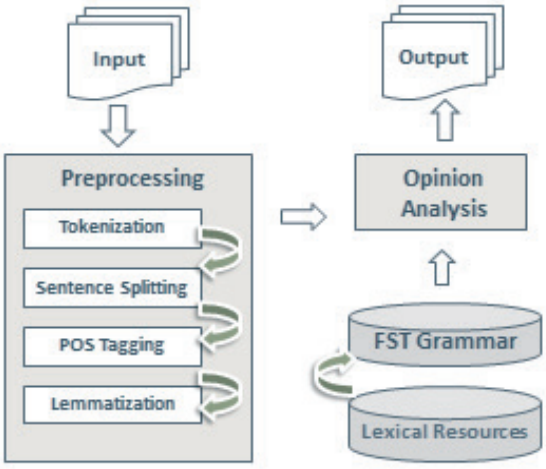

Figure 1 | Workflow for Opinion Mining

film noir, and social movies) manually collected from *athinorama.gr*. To test the method, we used unseen restaurant reviews (501 sentences) from *ask4food.gr*. In addition, the performance of the method was tested in a new domain, which was hotel reviews (475 sentences) manually selected from *trivago.gr*. Both test datasets were manually annotated with opinion expressions and opinion targets along with the respective sentiment polarity labels using GATE[1]. The annotation process resulted in 1969 annotations in total in both datasets. Table 2 provides more information about the annotations per test set.

	Restaurants	**Hotels**
Opinion Expressions	430	675
Opinion Targets	336	528
Total	766	1203

Table 2 | Number of annotations per test set

1 2]https://gate.ac.uk/sale/tao/splitch8.html

3.2. Lexical Resources

EvalLex (Pontiki et al. 2013) is an Appraisal Theory (Martin and White 2005) grounded Lexicon for Evaluative Language that was manually compiled for the Greek language. Originally, it contained 4429 evaluative terms and it was designed to enable a given Sentiment Analysis system to accurately detect "criticism" for negative evaluations and "praise" for positive evaluations. Each term was assigned a label according to its category (i.e. adjective (JJ), adverb (RB), or noun (NN)) and its sentiment polarity (i.e. negative (n), positive (p), or both (b)). In addition, adjectives and adverbs were further classified as follows based on the strength degree of their Evaluative Meaning (EM) and Prior Polarity (PP):

(1) Strong EM with a strong (p/n) PP e.g. "*υπεροπτικός" ("arrogant")* [JJ1n].
(2) Weak EM with a strong (p/n) PP e.g. "*ώριμος" ("mature")* [JJ2p].
(3) Strong or weak EM with a weak (p/n/b) or no PP e.g. "*μικρός" ("small")* [JJ3b].

In this work, EvalLex is expanded in two ways. First, nouns have been also manually classified into semantic categories according to the above three-way classification schema. Second, verbs and other types of sentiment vocabulary (e.g., informal and slang) have been added reaching a number of 5887 entries in total. In addition, we have manually compiled a lexicon of Modifiers that contains intensifiers (e.g., "*τελείως*" ("*totally*")), downtoners (e.g., "*κάπως*" ("*somewhat*")), and negators (e.g., "*καθόλου*" ("*not at all*")).

3.3. Opinion Grammar

The Opinion Grammar is the implementation of a four-phase algorithm where the output of each phase is input for the next one. Each phase comprises several modules (38 in total) that contain a variety of contextual lexico-syntactic patterns (340 in total). The patterns are templates that generate rules in the context around the candidate opinions and targets.

In a first phase (Phase 1), a set of rules perform stepwise contextual *Opinion Detection* (OD) and *Opinion Polarity Disambiguation* (OPD) using a) combinations be-

tween different types of EvalLex's JJs, RBS and NNs based on their EM and PP labels, and between EvalLex's entries, b) negation rules, and c) neutralization (pruning) rules that reject detected opinion expressions when they appear in particular contexts (e.g., conditional sentences). For example, pattern (1) below marks as opinions all the spotted sequences between EvalLex's adverbs and adjectives (e.g., "*εκνευριστικά μικρός*" *("annoyingly small")*) and assigns the respective polarity label taking into account the strength of the PP of each word (e.g. the opinion expression "*εκνευριστικά μικρός*" is assigned a negative polarity label):

(1) If Token[i] ∈ EvalLex & Token[i].Type ==“JJ3”
& Token[i-1] ∈ EvalLex & Token[i-1].Type == “RB[1|2]”
then Token[i].Label = “Sentiment”
& Token[i].Polarity == Token[i-1].Polarity

In the second phase (Phase 2), a set of rules perform OD, OPD and opinion target extraction (OTE) taking into account the output of Phase 1 and using in addition: a) lexical-level features (e.g. POS tags), b) candidate opinion target sequences before and after the detected opinions and, c) propagation rules (Qui et el. 2011). For example, pattern (2) marks as opinions and targets specific (sequences of) tokens when they appear in conjunction with already detected opinions:

(2) If Token[i] ∈ Sentiment & Token[i].Type ==“JJ”
& Token[i+1].POS==“CjCo” & Token[i+2] ∈ EvalLex
& Token[i+2].Type == “JJ” & Token[i+3].POSTag==“At”
& Token[i+4] ∈ Candidate_Target
then Token[i+2].Label = “Sentiment”
& Token[i+2].Polarity ==Token[i].Polarity
& Token[i+4].Label = “target”
& Token[i+4].Polarity ==Token[i].Polarity

Again, the output of Phase 2 is the input for Phase 3, where OD, OPD and OTE is performed based in addition on a) EvalLex's VBs, and b) macros of trigger verbs (e.g., "*θεωρώ*" ("*believe*"), "*αποτελώ*" ("*constitute*")). Finally, Phase 4 focuses only on OTE using rules that impose less constrains to the context before and after detected opinions in order to capture longer dependencies.

The method is precision-oriented; the goal is to incrementally increase the recall in each phase while maintaining high precision. In addition, the proposed method is domain-independent in that it does not make use of any domain specific knowledge (e.g., domain specific sentiment vocabulary, predefined lists of aspect terms i.e. terms that name particular aspects of a target entity).

4. Evaluation

The performance of the opinion analyzer is measured on the test datasets in terms of Precision (P), Recall (R) and F-Measure (F-1) using the GATE Annotation Diff Tool. P, R and F-1 are defined as follows:

$$P = \frac{|S \cap G|}{|S|} \qquad R = \frac{|S \cap G|}{|G|} \qquad F_1 = \frac{2 \cdot P \cdot R}{P + R}$$

S is the set of the opinion expressions and opinion targets that the system returned for all the test texts (of each domain), and *G* is the set of the gold (correct) opinion expressions and target annotations. F1 score is the harmonic mean of P and R.

	Restaurants			**Hotels**		
	Precision	**Recall**	**F-Measure**	**Precision**	**Recall**	**F-Measure**
Opinion detection (OD)	88%	66%	76%	93%	67%	78%
Opinion polarity disambiguation (OPD)	80%	60%	69%	84%	60%	80%
Opinion target extraction (OTE)	74%	42%	54%	83%	48%	74%

Table 2 | Evaluation Results

The evaluation results confirmed our expectations favoring a precision-orientated and domain-independent method, since it achieves high precision in all types of the extracted information in both test domains. The performance in opinion polarity disambiguation is lower mainly due to the fact that, currently, the opinion analyzer performs only binary (positive/negative) sentiment classification not taking into account the neutral (i.e. mildly positive or mildly negative) opinions. The performance is better in the unknown (hotels) test domain as compared to the restaurants one due to the fact that the latter contains less opinions (annotations), which, in addition are about a particular restaurant's degustation menu with complex dish names (e.g., "*cheesecake με καραμελωμένο ανανά και σάλτσα μάγκο*"). As expected opinion target extraction is a harder and more complex task as compared to opinion detection. Overall, the performance is negatively affected mainly due to a) the limitations of the shallow syntactic relations modelling, since long distance dependencies cannot be captured through a window of a limited number of tokens, and b) cases that the method is not designed to address yet (i.e. co-reference resolution, domain specific sentiment vocabulary, complex linguistic phenomena like metaphors and irony).

Figure 2 | *Summary of opinion expressions extracted from the restaurants test dataset*

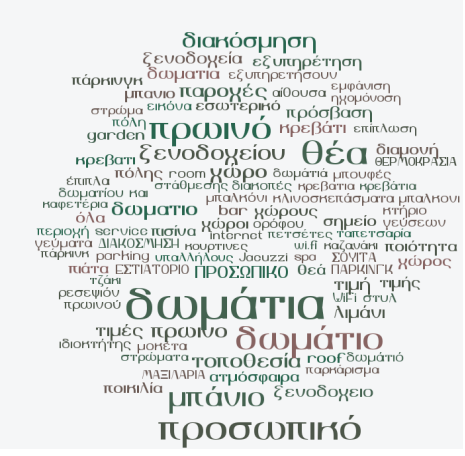

Figure 3 | *Summary of opinion targets extracted from the hotels test dataset*

5. Conclusions

In this paper we presented a rule-based method for the automatic detection of opinions and their targets in Greek review texts. The evaluation results confirmed our expectations favoring a precision-orientated and domain-independent method, since it achieves substantial precision in all types of the extracted information in both test domains. The output of the opinion analyzer can be used to generate insightful sentiment summaries (e.g. Figures 2 and 3). In addition, the method can serve as a core component for aspect based sentiment analysis, namely map the extracted opinion targets to aspect categories (e.g. map opinion targets like "fish", "pasta", and "salad" to the coarse predefined category "food" in the restaurants domain), and/or enhance the grammar with domain specific knowledge (i.e. lists of aspect terms as candidate opinion targets). Future work involves testing on more domains and experimenting with other types of techniques to deal with the limitations of the current method and improve the recall.

Acknowledgments

Research presented in this paper is supported by the POLYTROPON (KRIPIS-GSRT, MIS: 448306) project.

References

Agathagelou, Pantelis, Ioannis, Katakis, Fotios, Kokkoras, and Konstantinos Ntonas. 2014. "Mining Domain-specific Dictionaries of Opinion Words". In *Proceedings of the 15th International Conference* on *Web Information System Engineering* (*WISE 2014*), *Thessaloniki, Greece*, 47–62.

Agrawal, Rakesh, and Ramakrishnan Srikant. 1994. "Fast Algorithms for Mining Association Rules in Large Databases". In *Proceedings of the 20th International Conference on Very Large Data Bases*, Santiago de Chile, Chile, 487–499.

Carenini, Giuseppe, Raymond T., Ng, and Ed Zwart. 2005. "Extracting Knowledge from Evaluative Text". In *Proceedings of the 3rd International Conference on Knowledge Capture*, Banff, Alberta, Canada, 11–18.

Cunningham, Hamish, Diana, Maynard, and Valentin Tablan. 2000. "JAPE: a Java Annotation Patterns Engine". *Research Memorandum CS–00–10*, Department of Computer Science, University of Sheffield.

Ding, Xiaowen, Liu, Bing, and Philip S. Yu. 2008. "A Holistic Lexicon-based Approach to Opinion Mining". In *Proceedings of 1st ACM International Conference on Web Search and Data Mining (WSDM-2008)*, Stanford, California, USA, 231–240.

Duan, Wenjing, Bin, Gu, and Andrew B. Whinston. 2008. "The Dynamics of Online Word-of-Mouth and Product Sales—An Empirical Investigation of the Movie Industry". *Journal of Retailing* 84(2):233–242.

Hu, Mingqing, and Liu Bing. 2004. "Mining and Summarizing Customer Reviews". In *Proceedings of 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Seattle, WA, USA, 168–177.

Jakob, Niklas, 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 (EMNLP 2010)*, Massachusetts, USA, 1035–1045.

Karampatsis Rafael Michael. 2014. "Social Media Sentiment Analysis". M.Sc. Thesis. Athens University of Economics and Business.

Lafferty, John, Andrew, McCallum, and Fernando Pereira. 2001. "Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data". In *Proceedings of the Eighteenth International Conference on Machine Learning (ICML'01)*, San Francisco, CA, USA, 282–289.

Li, Fangtao, Chao, Han, Minlie, Huang, Xiaoyan, Zhu, Ying-Ju Xia, Shu, Zhang, and Hao Yu. 2010. "Structure-aware Review Mining and Summarization". In *Proceedings of the 23rd International Conference on Computational Linguistics (COLING-2010)*, Beijing, China, 653–661.

Liu, Bing, and Lei Zhang. 2013. "A Survey of Opinion Mining and Sentiment Analysis". *Mining Text Data,* 415–463.

Liu, Kang, Xu, Liheng, Liu, Yang, and Zhao Jun. 2013. "Opinion Target Extraction using Partially Supervised Word Alignment Model". In *Pro-*

ceedings of the International Joint Conference on Artificial Intelligence, Beijing, China, 2134–2140.

Liu, Kang, Xu, Linheng, and Zhao Jun. 2015. "Co-extracting Opinion Targets and Opinion Words from Online Reviews based on the Word Alignment Model". *IEEE Transactions on Knowledge and Data Engineering* 27(3):636–650.

Ma, Tengfei, and Wan Xiaojun. 2010. "Opinion Target Extraction in Chinese News Comments". In *Proceedings of 23rd International Conference on Computational Linguistics (COLING-2010)*, Beijing, China, Poster Volume, 782–790.

Martin, James R., and Peter R. R. *White*. 2005. "The Language of *Evaluation*, Appraisal in English". Palgrave Macmillan, London & New York.

Papageorgiou, Haris, Prokopis, Prokopidis, Iason, Demiros, Voula, Giouli, Alexis, Konstantinidis, and Stelios Piperidis. 2002. "Multi–level XML–based Corpus Annotation." In *Proceedings of the 3rd International Conference on Language Resources and Evaluation (LREC 2002)*, Las Palmas, Spain, 1723–1728.

Pang, Bo, and Lillian Lee. 2008. "Opinion Mining and Sentiment Analysis". *Foundations and Trends in Information Retrieval* 2(1-2):1–135.

Pontiki, Maria, Zoe, Angelou, and Haris Papageorgiou. 2013. "Sentiment Analysis: Building Bilingual Lexical Resources". In *Proceedings of the 11th International Conference on Greek Linguistics*, Rhodes Island, Greece, 1390–1400.

Pontiki, Maria, Dimitrios, Galanis, John, Pavlopoulos, Haris, Papageorgiou, Ion, Androutsopoulos, and Suresh Manandhar. 2014. "SemEval-2014 Task 4: Aspect Based Sentiment Analysis". *In Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, Dublin, Ireland, 27–35.

Pontiki, Maria, Dimitrios, Galanis, Haris, Papageorgiou, Suresh, Manandhar, and Ion Androutsopoulos. 2015. "SemEval-2015 Task 12: Aspect Based Sentiment Analysis". In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, Denver, Colorado, 486–495.

Popescu, Ana-Maria, and Oren Etzioni. 2005. "Extracting Product Features and Opinions from Reviews". In Proceedings of the 2005 Conference on Human Language Technology and Empirical Methods in Natural Language Processing, Vancouver, British Columbia, Canada, 339–346.

Prokopidis, Prokopis, Byron, Georgantopoulos, and Haris Papageorgiou. 2011. "A Suite of NLP Tools for Greek". In *Proceedings of the 10th Inter-*

national Conference of Greek Linguistics, Komotini, Greece, 373–383.

Qui, Guang, Liu, Bing, Bu, Jianjun, and Chen Chun. 2011. "Opinion Word Expansion and Target Extraction through Double Propagation". *Computational Linguistics* 37(1):9–27.

Solakidis, Georgios, Konstantinos, Vavliakis, and Pericles Mitkas. 2014. "Multilingual Sentiment Analysis using Emoticons and Keywords". In *Proceedings of the 2014 IEEE/WIC/ACM International Joint Conferences on Web*, Warsaw, Poland, 102–109.

Wu, Yuanbin, Qi, Zhang, Xuanjing, Huang, and Lide Wu. 2009. "Phrase Dependency Parsing for Opinion Mining." In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing (EMNLP 2009)*, Singapore, 1533–1541.

Yoo, Kyung Hyan, and Ulrike Gretzel. 2008. "What Motivates Consumers to Write Online Travel Reviews?" *Information Technology & Tourism* 10:283–295.

Zhang, Qi, Yuanbin, Wu, Tao Li, Mitsunori, Ogihara, Joseph, Johnson, and Xuanjing Huang. 2009. "Mining Product Reviews based on Shallow Dependency Parsing". In *Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval (SIGIR '09)*, New York, USA, 726–727.