Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence

Opinion Target Extraction Using a Shallow Semantic Parsing Framework

Shoushan Li, Rongyang Wang, Guodong Zhou*

Natural Language Processing Lab, School of Computer Science and Technology
Soochow University, Suzhou, 215006, China
{shoushan.li, wangrongyang.nlp}@gmail.com, gdzhou@suda.edu.cn

Abstract

In this paper, we present a simplified shallow semantic parsing approach to extracting opinion targets. This is done by formulating opinion target extraction (OTE) as a shallow semantic parsing problem with the opinion expression as the predicate and the corresponding targets as its arguments. In principle, our parsing approach to OTE differs from the state of the art sequence labeling one in two aspects. First, we model OTE from parse tree level, where abundant structured syntactic information is available for use, instead of word sequence level, where only lexical information is available. Second, we focus on determining whether a constituent, rather than a word, is an opinion target or not, via a simplified shallow semantic parsing framework. Evaluation on two datasets shows that structured syntactic information plays a critical role in capturing the domination relationship between an opinion expression and its targets. It also shows that our parsing approach much outperforms the state of the art sequence labeling one.

1. Introduction

Recent years have witnessed an exploding interest in sentiment analysis in natural language processing and data mining due to its inherent challenges and wide applications. One fundamental problem in sentiment analysis is opinion target extraction (OTE) which aims to identify topics on which an opinion is expressed (Pang and Lee, 2008). For example, in product reviews, opinion targets are often the product itself (e.g. "*I absolutely love this product.*") or its specific features, such as design and quality (e.g. "*The design of iphone 4S is fantastic.*", "*They are of very high quality.*"). Previous approaches to this task mainly focus on unsupervised learning where some heuristic rules are usually designed to recognize the opinion targets (Hu and Liu, 2004). Basically, designing the heuristic rules is normally difficult and often suffers from low performance. More recently, supervised learning approaches to OTE have attracted an increasing interest. Although supervised learning approaches normally much outperform unsupervised learning approaches to OTE with the help of annotated data (Zhuang et al., 2006), their performances are normally far from expectation and remain challenging due to following critical issues.

First, although OTE can be considered as a specific task of information extraction (IE) (Cowie and Lehnert, 1996), the concerned information here (i.e. opinion targets) is highly bound to an opinion expression, different from many traditional IE tasks. Correctly extracting opinion targets needs not only to consider the context of the targets themselves but also to determine whether the targets are related to an opinion expression or not. It is still a difficult issue to model the close relationship between an opinion expression and its targets in supervised learning approaches.

Second, OTE is a fine-grained task. Different from those coarse-grained ones like document-level sentiment classification (Pang et al., 2002), simply employing word tokens or part-of-speech features no longer qualifies for successful OTE. In contrast, deep knowledge, such as the sentence-level syntactic structure, becomes essential to successful OTE. In fact, several unsupervised approaches have noticed this challenge and employed syntactic knowledge, e.g. syntactic patterns and dependency relationship, to this task with some performance improvement (Kobayashi et al., 2007; Qiu et al., 2011). However, there is only a few attempts on how to employ syntactic knowledge in supervised approaches to OTE (Kim et al., 2008).

In this paper, we explore supervised OTE from a parse tree structure perspective and formulate it as a shallow semantic parsing problem, which has been extensively studied in the past few years (Xue, 2008). In particular, the opinion expression is recast as the predicate and the corresponding opinion targets are recast as its arguments. The motivation behind is that (1) the parse tree structure includes various paths from the opinion expression to the opinion targets, which naturally provide a reasonable way to capture the close relationship between the opinion targets and the opinion expression, so as to handle the first challenge; (2) the parse tree structure provides abundant syntactic knowledge

---

* Corresponding author

to better recognize opinion targets, so as to handle the second challenge. In principle, recasting an opinion target as a constituent in a parse tree provides more potential to better represent its close relationship with the opinion expression than as a string in a word sequence.

Our parsing approach to supervised OTE differs from existing studies in two aspects. First, we extend OTE from the word sequence level into the parse tree level, where structured syntactic information is available. Second, we focus on determining whether a constituent in a parse tree, rather than a string in a word sequence, is an opinion target or not. Evaluation on two datasets shows that our parsing approach much outperforms the state-of-the-art sequence labeling one by Jakob and Gurevych (2010).

2. Related Work

While there is a certain mount of literature within the NLP community on unsupervised OTE (Hu and Liu, 2004; Popescu and Etzioni, 2005; Blei and Jordan, 2006; Bloom et al., 2007; Kim and Hovy, 2006; Titov and McDonald, 2008), supervised learning to OTE is relatively new.

Zhuang et al. (2006) obtain various dependency relationship templates from an annotated movie corpus and apply them to supervised OTE. Empirical evaluation shows that their template-based classification approach greatly outperforms the unsupervised one by Hu and Liu (2004).

Kessler and Nicolov (2009) model OTE as a ranking problem and extract the highest ranked candidates as opinion targets. Empirical evaluation shows that their candidate ranking approach outperforms several unsupervised ones.

Jakob and Gurevych (2010) model OTE as a word sequence labeling problem. Empirical evaluation shows that their sequence labeling approach much outperforms both the template-based classification approach by Zhuang et al. (2006) and the candidate ranking approach by Kessler and Nicolov (2009), representing the state-of-the-art in supervised OTE.

Instead, our constituent parsing approach addresses OTE from a parse tree structure perspective.

3. DSRC Corpus

This study employs the DSRC corpus[1], as described in Toprak et al. (2010), which contains two datasets: university and web-service. In the corpus, every sentence is annotated with opinion expressions, and their corresponding opinion holders and opinion targets, and so on. Table 1 shows an example of such annotation scheme. In this paper, we only focus on the opinion targets.

Table 1: Example of an annotated sentence in the DSRC corpus

Data	<word id "word 1">I've</word> <word id "word 2">always</word> <word id "word 3">been</word> <word id "word 4">pretty</word> <word id "word 5">dubious</word> <word id "word 6">about</word> <word id "word 7">the</word> <word id "word 8">concept</word> <word id "word 9">of</word> <word id "word 10">online</word> <word id "word 11">universities</word>
Markables	1. <markable span "word 10..word11" **annotation_type= "target"** /> 2. <markable span "word 5" **annotation_type="opinion expression"** />

Table 2 gives the statistics of each dataset. From this table, we can see that the average length of opinion targets is less than two, with 86.81%/8.17%/4.17%/0.85% containing one/two/three//more words in the university dataset and 81.46%/12.84%/5.37%/0.32% containing one/two/three/more words in the web-service dataset.

Table 2: Statistics of the DSRC corpus

Number	University	Web services
Documents	256	234
Sentences	2911	7575
Sentences with opinion	**1012**	**1372**
Targets	1175	1861
Target types	335	661
Average length of targets	1.48	1.37

For preprocessing, all the sentences in the DSRC corpus are parsed using the Stanford Parser[2], which is a Java implementation of probabilistic natural language parsers, including both a highly optimized PCFG parser and a lexicalized dependency parser (Klein and Manning, 2003).

4. OTE via Shallow Semantic Parsing

In this section, we first formulate the OTE task as a shallow semantic parsing problem. Then, we deal with it using a simplified shallow semantic parsing framework.

4.1 Formulating OTE as a Shallow Semantic Parsing Problem

Given a parse tree and a predicate in it, shallow semantic parsing recognizes and maps all the constituents in the sentence into their corresponding semantic arguments (roles) of

[1] http://www.ukp.tu darmstadt.de/data/sentiment analysis/

[2] http://nlp.stanford.edu/software/lex parser.shtml#Citing

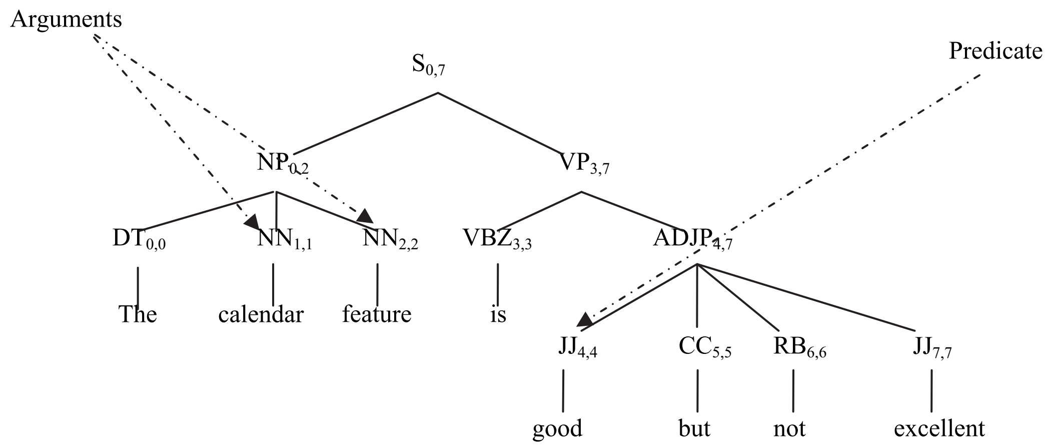

Figure 1: Illustration of an opinion expression (predicate) and its corresponding opinion targets (arguments) in a parse tree

the predicate. As far as OTE considered, the opinion expression can be regarded as the predicate, while the opinion targets can be mapped into its arguments. For example, in the sentence "*The calendar feature is good but not excellent.*" as shown in Figure 1, two opinion expressions of $JJ_{4,4}$ and $JJ_{7,7}$ are found and the opinion target consists of two words: $NN_{1,1}$ and $NN_{2,2}$. In this study, we assume opinion expressions have been recognized and treat the nearest opinion expression to an opinion target as its predicate. Thus, in this example, $JJ_{4,4}$ is considered as a predicate and $NN_{1,1}$ and $NN_{2,2}$ are considered as two arguments.

In particular, given a opinion expression and one of its opinion targets, which contains *m* words: $word_1$, ..., $word_m$, we adopt following three heuristic rules to map the opinion target into several constituents which can be deemed as its arguments in the given parse tree.

(1) The opinion expression itself and all of its ancestral constituents are non-arguments.

(2) If all child constituents of constituent *X* are recognized as arguments, then *X* is labeled as an argument and all its child constituents are re-labeled as non-arguments.

(3) If not all of the child constituents of constituent *X* are recognized as arguments, then *X* is labeled as a non-argument.

The first rule ensures that no argument covers the opinion expression while the remaining two rules ensure no overlap between any two arguments in an opinion target. These constraints between a predicate and its arguments are somehow consistent with shallow semantic parsing (Carreras and Màrquez, 2005). For example, $NN_{1,1}$, $NN_{2,2}$ and $NP_{0,2}$ cannot be arguments at the same time because $NP_{0,2}$'s child $DT_{0,0}$ is not a argument ($NN_{1,1}$, $NN_{2,2}$ and $NP_{0,2}$ are overlapped).

Compared with traditional shallow semantic parsing which needs to assign an argument with a semantic label, OTE does not involve semantic label classification and thus can be simplified into three phases: argument pruning, argument identification and post-processing.

4.2 Argument Pruning

Argument pruning aims to filter out those constituents which are most likely non-arguments of a predicate according to some heuristic rules. Here, we propose two pruning rules as follows:

(1) The predicate constituent itself and its ancestral constituents in the parse tree are filtered out as non-arguments.

(2) The constituents which contain more than three leaf nodes in the parse tree are filtered out as non-arguments and instead, their child constituents are considered as arguments candidates individually.

Here, the first rule is adopted mainly due to the first constraint as described in Section 4.1 while the second rule is proposed mainly due to the statistics of opinion targets. Generally, an opinion target contains less than four words. For example, as pointed in Section 3, only 0.85% of opinion targets contain more than three words in the DSRC corpus. Furthermore, we can simply merge the separated arguments to form an opinion target when it contains more than three words. In this way, many of non-argument constituents can be filtered out safely and conveniently. Take Figure 1 as an example, $S_{0,7}$, $VP_{3,7}$, and $ADJP_{4,7}$ are filtered out according to the first rule since they are ancestral constituents of the predicate constituent $JJ_{4,4}$.

4.3 Argument Identification

For remaining argument candidates, we employ a binary classifier to determine whether an argument candidate is an argument or not, using following two groups of features: basic features and additional features.

Basic Features

Table 3 lists the basic features for argument identification. These features are directly related with the predicate and the argument candidate, and have been widely used in common shallow semantic parsing for both verbal and nominal predicates (Xue, 2008).

Table 3: Basic features and their instantiations for OTE, with $NN_{1,1}$ (*calendar*) as the focus argument candidate and $JJ_{4,4}$ (*good*) as the given predicate, with regard to Figure 1.

Feature	Remarks
B1	The opinion expression (*good*)
B2	The syntactic category of the argument candidate (*NN*)
B3	The headword of the argument candidate (*calendar*)
B4	The POS of the headword of the argument candidate (*NN*)

Additional Features

To capture more useful information in opinion targets and opinion expressions, we also explore various kinds of additional features in capturing more details regarding the argument candidate and the predicate, as shown in Table 4. In particular, we categorize the additional features into three groups according to their relationship with the argument candidate (Arg, in short) and the given predicate (Pre, in short). In particular, various parsing paths are included to capture the relationship between the opinion target and the opinion expression.

Since some proposed features may not be effective in argument identification, we adopt a greedy feature selection algorithm, as described in Jiang and Ng (2006), to pick up effective features incrementally according to their contributions on the development data. Specially, the algorithm repeatedly selects one feature each time which contributes most, and stops when adding any of the remaining features fails to improve the performance. As far as OTE concerned, the whole feature selection process could be done by first running the selection algorithm with the basic features (B1-B4) and then incrementally picking up effective features from the additional features (Arg1 - Arg2, Pre1 - Pre2, and A-P1 - A-P8).

4.4 Post-Processing

As mentioned in Section 3, 86.81%/8.17%/4.17%/0.85% of the targets contain one/two/three/more words respectively. However, our opinion target extractor may return more long words than expected. We also note that, in the original corpus, most targets exclude starter determiners and pronouns, such as "*a*", "*an*", "*the*" and "*this*". For example, in the sentence of "*The calendar feature is good but not excellent.*", determiner "*the*" is not included in opinion target "*calendar feature*" in the annotation. However, our opinion target extractor recognizes the whole noun phrase "*the calendar feature*" as an opinion target. In fact, the annotation on this type of opinion targets is not consistent in the original corpus: some starter determiners and pronouns are contained in the opinion targets while others not. To keep consistency, we delete such starter determiners and pronouns from opinion targets in the annotated corpus. Overall, 0.55% of the opinion targets are concerned. Similarly, during post-processing, we delete those starter determiners and pronouns from extracted opinion targets.

Table 4: Additional features and their instantiations for OTE, with $NN_{1,1}$ (*calendar*) as the focus argument candidate and $JJ_{4,4}$ (*good*) as the given predicate, with regard to Figure 1.

Feature	Remarks
Argument candidate (Arg) related context features	
Arg1	The left word and its POS. (*the*, *DT*)
Arg2	The right word and its POS. (*feature*, *NN*)
Predicate (Pre) related context features	
Pre1	The left word and its POS. (*is*, *VBZ*)
Pre2	The right word and its POS. (*but*, *CC*)
Arg-Pre-related structured features	
A-P1	The subcategory governing the predicate and the argument candidate (*NP:DT+NN+NN*)
A-P2	The syntactic path from the argument candidate to the predicate. (*NN>NP>S<VP<ADJP<JJ*)
A-P3	The number of the nodes in A-P2 (*6*)
A-P4	Compressed A-P2: compressing sequences of identical labels into one. (*NN>S<VP<ADJP<JJ*)
A-P5	The syntactic partial path from the argument candidate to the least governing node of both the argument candidate and the predicate (*NN>NP>S*)
A-P6	The syntactic partial path from the predicate to the least governing node of both the argument candidate and the predicate (*S<VP<ADJP<JJ*)
A-P7	Whether there is a clause tag (*S*) between the predicate and the argument candidate (*yes*)
A-P8	The positional relationship of the argument candidate with the predicate:. "left" or "right". (*left*)

5. Experiments

In this section, we will systematically evaluate our simplified shallow semantic parsing approach to opinion target extraction on the DSRC corpus.

5.1 Experimental Setting

Dataset: The DSRC corpus is used to evaluate our approach. For details, please refer to Section 3. Same as Jakob and Gurevych (2010), we assume that the opinion expressions in the opinion sentences are known and only the opinion sentences are used for evaluation.

Classification algorithm: Standard classification algorithms, such as Support Vector Machine and Maximum Entropy, can all be employed. In order to fairly compare our approach to the-state-of-the-art sequence labeling one by Jakob and Gurevych (2010), which adopt CRF in sequence labeling, we also use CRF, however as a classifier. That is, our implementation is essentially different from the sequence labeling approach. In our approach, we adopt CRF as a classifier with constituent as the basic classification unit, e.g., "*the calendar feature*", to determine whether a constituent is an argument or not. In comparison, Jakob and Gurevych (2010) adopt CRF as a sequence labeling tool with word as the basic labeling unit to determine whether a word is at the beginning, at the non-beginning and outside of an argument.

Evaluation metrics: Exact match is used to evaluate the correctness of an extracted opinion target. That is to say, an extracted opinion target is considered as correct only if it has exactly the same span boundaries as the annotated ones in the gold standard. Same as Jakob and Gurevych (2010), the precision (P), recall (R) and F-measure (F) of the extracted opinion targets are employed as evaluation metrics.

5.2 Contribution of Different Features

In this experiment, all the opinion sentences in each dataset are randomly divided into 90% and 10% which are used as training data and development data respectively.

Contribution of Basic Features on Development Data

Table 5 lists the performance of basic features on the development data. It shows that the basic features achieve the performance of 56.72 and 56.95 in F1-measure on the university and web-service domains respectively. In particular, the headword of the argument candidate (B3) contributes most individually and achieves 50.65 and 50.10 in F-measure on the two datasets respectively.

Table 5: Contribution of basic features

	Domain: University			Domain: Web Service		
	P(%)	*R*(%)	*F*	*P*(%)	*R*(%)	*F*
B1	40.33	2.79	5.16	38.59	3.24	5.91
B2	52.71	13.40	21.07	62.83	27.57	38.23
B3	74.65	38.82	50.65	73.76	38.07	50.10
B4	59.59	28.10	37.75	67.34	37.87	48.38
ALL	69.82	48.07	56.72	69.35	48.53	56.95

Contribution of Additional Features on Development Data

Tables 6-8 show the contribution of each category of additional features respectively with all the four basic features already included. From these tables, we can see that except feature Pre1, all the Arg-related and Pre-related context features are of little help, while except A-P1 in the web-service domain, all the Arg-Pre-related structured features are very effective, especially those path features generated from the parse tree. This verifies the importance of employing structured syntactic knowledge for OTE.

Table 6: Contribution of additional argument candidate related context features

	Domain: University			Domain: Web Service		
	P(%)	*R*(%)	*F*	*P*(%)	*R*(%)	*F*
Basic	69.82	48.07	56.72	69.35	48.53	56.95
Arg1	68.9	47.8	56.17	69.17	47.84	56.42
Arg2	69.8	48.02	56.62	69.37	48.31	56.80

Table 7: Contribution of additional predicate related context features

	Domain: University			Domain: Web Service		
	P(%)	*R*(%)	*F*	*P*(%)	*R*(%)	*F*
Basic	69.82	48.07	56.72	69.35	48.53	56.95
Pre1	69.45	50.01	57.89	68.35	49.54	57.37
Pre2	67.56	49.32	56.67	67.63	48.56	56.39

Table 8: Contribution of additional Arg-Pre related structured features

	Domain: University			Domain: Web Service		
	P(%)	*R*(%)	*F*	*P*(%)	*R*(%)	*F*
Basic	69.82	48.07	56.72	69.35	48.53	56.95
A P1	70.51	49.04	57.57	69.38	47.91	56.54
A P2	73.96	50.32	59.62	73.19	50.9	59.89
A P3	73.01	54.29	61.92	72.62	54.35	62.07
A P4	74.39	51.3	60.53	74.06	52.31	61.14
A P5	74.09	52.5	61.19	73.59	54.95	62.81
A P6	73.65	54.13	62.08	72.24	51.7	60.12
A P7	72.01	52.84	60.66	70.45	49.98	58.36
A P8	70.52	49.57	57.96	69.16	49.2	57.33

Moreover, we perform the greedy algorithm as described in Section 4.2 to select a set of optional additional features on the development data of the university dataset. Table 9 shows the effect of selected features in an incremental way. It shows that using the additional feature set of {A-P2, A-P4, A-P5, Pre2, A-P6, A-P8, A-P7} achieves the best performance and significantly improves the performance by 9.56 in F-measure from 56.00 to 65.56 (p-value<0.05). We don't include other additional features since their contributions can be largely ignored. This also applies to the web-service dataset. Therefore, we only include those effective additional features as shown in Table 9 plus all the four basic features in our remaining experiments.

Table 9: Performance improvement of including additional features in an incremental way on the university dataset

	Domain: University		
	P(%)	*R*(%)	*F*
Basic	64.90	49.25	56.00
+A P2	70.39	53.77	60.97
+A P4	75.00	54.27	62.97
+A P5	71.15	55.78	62.54
+Pre2	71.25	57.29	63.51
+A P6	71.43	57.79	63.89
+A P8	72.05	58.29	64.44
+A P7	73.29	59.30	65.56

5.3 Comparison with the State-of-the-Art

For comparison, we re-implement the state-of-the-art sequence labeling approach for OTE, as proposed by Jakob and Gurevych (J and G, in short) (2010), who employ CRF with various features, such as token, POS, short dependency path, word distance, and opinion sentence characteristics. Specially, the dependency path is obtained using the Stanford Parser. For better comparison, all the opinion sentences in each dataset are randomly divided into 10 folds so as to perform 10-fold cross-validation. The reported results are the average ones over the 10 runs. In our approach, we use the best feature set obtained above. Besides, the comparison is conducted in both in-domain and cross-domain ways.

In-domain comparison

Figure 2 compares the performances of our parsing approach and the state-of-the-art sequence labeling one by J and G (2010). It shows that our parsing approach significantly outperforms J and G (2010) in both datasets (*p*-value<0.05 in F-measure). Our parsing approach is especially advantageous in terms of recall. That is, we observe that our parsing approach is capable of finding more opinion targets which are ignored by J and G (2010). For example, in the sentence of "*The folks in the financial aid department were also extremely helpful*", the sequence labeling approach fails to find "*folks*" as an opinion target since opinion target “folks” is far apart from opinion expression "helpful" while our parsing approach "succeeds", due to the effectiveness of structured syntactic information in our simplified shallow semantic approach. Moreover, the sequence labeling approach sometimes can only extract partial opinion targets while our approach is more effective due to its considering a constitute as an opinion target instead of a string in a word sequence.

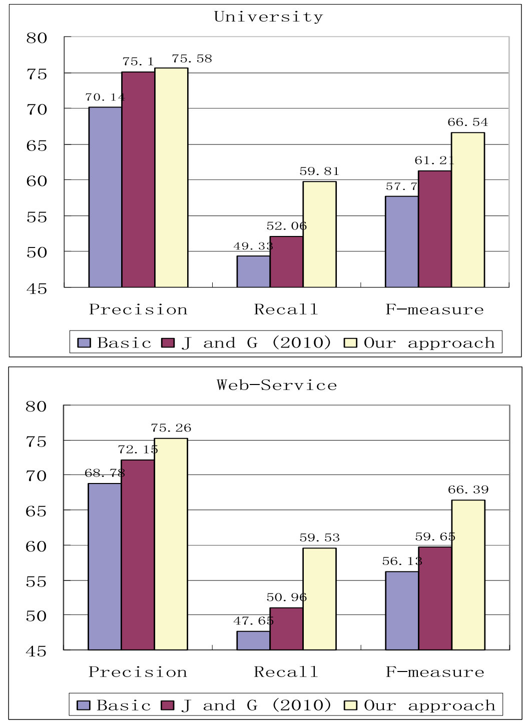

Figure 2: In domain performance comparison of our parsing approach with the state of the art sequence labeling approach by J and G (2010)

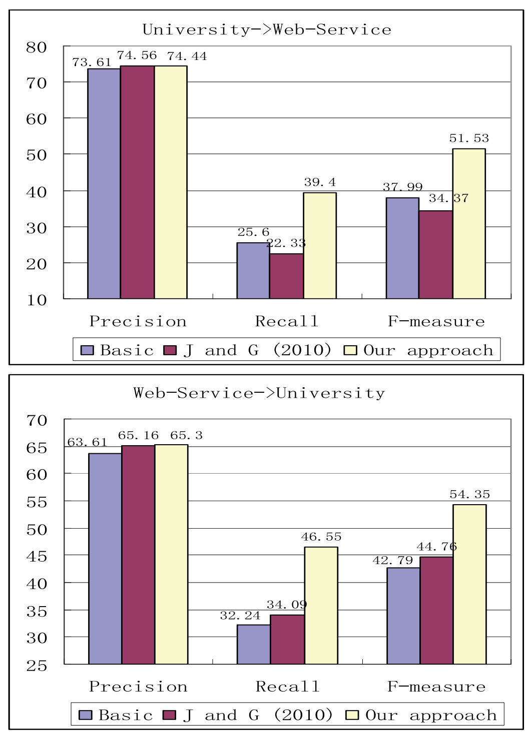

Figure 3: Cross domain performance comparison of our pars ing approach with the state of the art sequence labeling approach by J and G (2010)

Cross-domain comparison

In this setting, all the opinion sentences in each dataset are used for training an OTE system, which is then used to test on another dataset (a different domain). Figure 2 compares the performances of our parsing approach and the state-of-the-art sequence labeling approach by J and G (2010). It shows that our parsing approach significantly outperforms J and G (2010) in both cross-domain evaluations (*p*-value<0.05) by 17.16 in 9.59 in F-measure. The excellent performance on cross-domain OTE once again verifies the importance of structured syntactic information in the OTE task.

6. Conclusion

In this paper, we present a novel approach to opinion target extraction by formulating it as a shallow semantic parsing problem. Experimental studies demonstrate that structured syntactic information plays a critical role in capturing the domination relationship between an opinion expression and an opinion target, and that our parsing approach significantly outperforms the state-of-the-art sequence labeling approach by Jakob and Gurevych (2010) in both in-domain and cross-domain settings.

For the future work, we will explore other types of information to opinion target extraction and include opinion holder extraction in our parsing framework. Moreover, we will systematically explore our parsing approach to cross-domain opinion target extraction and opinion holder extraction.

Acknowledgments

The research work described in this paper has been partially supported by three NSFC grants, No. 61003155, No. 60873150 and No. 90920004, one National High-tech Research and Development Program of China No. 2012AA011102 and Open Projects Program of National Laboratory of Pattern Recognition. We also thank the three anonymous reviewers for their helpful comments.

References

Blei D., A. Ng, and M. Jordan. 2006. Correlated Topic Models. In *Proceedings of NIPS 06*, 147 154.

Bloom K., N. Garg, and S. Argamon. 2007. Extracting Appraisal Expressions. In *Proceedings of HLT NAACL 07*, 308 315.

Carreras X. and L. Màrquez. 2005. Introduction to the CoNLL 2005 Shared Task: Semantic Role Labeling. In *Proceedings of CoNLL 05*, 152 164.

Cowie J. and W. Lehnert. 1996. Information Extraction. *Communications of the ACM*, 39(1):80 91.

Hu M. and B. Liu. 2004. Mining and Summarizing Customer Reviews. In *Proceedings of SIGKDD 04*, 168 177.

Jakob N. and I. Gurevych. 2010. Extracting Opinion Targets in a Single and Cross Domain Setting with Conditional Random Fields. In *Proceedings of EMNLP 10*, 1035 1045.

Jiang Z. and H. Ng. 2006. Semantic Role Labeling of NomBank: A Maximum Entropy Approach. In *Proceedings of EMNLP 06*, 138 145.

Kessler J. and N. Nicolov. 2009. Targeting Sentiment Expressions through Supervised Ranking of Linguistic Configurations. In *Proceedings of the Third International AAAI Conference on Weblogs and Social Media*, 90 97.

Kim S. and E. Hovy. 2006. Extracting Opinions, Opinion Holders, and Topics Expressed in Online News Media Text. In *Proceedings of the ACL Workshop on Sentiment and Subjectivity in Text*, 1 8.

Kim Y., S. Kim, and S. Myaeng. 2008. Extracting Topic related Opinions and their Targets in NTCIR 7. In *Proceeding of NTCIR 7 Workshop*, 247 254.

Klein D. and C. Manning. 2003. Accurate Unlexicalized Parsing. In *Proceedings of ACL 03*, 423 430.

Kobayashi, Nozomi, Kentaro Inui, and Yuji Matsumoto. 2007. Extracting Aspect evaluation and Aspect of Relations in Opinion Mining. In *Proceedings of EMNLP 07*, 1065 1074.

Pang B. and L. Lee. 2008. Opinion Mining and Sentiment Analysis: Foundations and Trends. *Information Retrieval*, vol.2(12), pp.1 135.

Pang B., L. Lee, and S. Vaithyanathan. 2002. Thumbs up? Sentiment Classification using Machine Learning Techniques. In *Proceedings of EMNLP 02*, pp.79 86.

Popescu A. and O. Etzioni. 2005. Extracting Product Features and Opinions from Reviews. In *Proceedings of EMNLP 05.* 339 346.

Qiu G., B. Liu, J. Bu, and C. Chen. 2011. Opinion Word Expansion and Target Extraction through Double Propagation. *Computational Linguistics*, vol.37(1), 10 27.

Titov I. and R. McDonald. 2008. Modeling Online Reviews with Multi grain Topic Models. In *Proceedings of WWW 08.* 111 120.

Toprak C., N. Jakob, and I. Gurevych. 2010. Sentence and Expression Level Annotation of Opinions in User Generated Discourse. In *Proceedings of ACL 10*, 575 584.

Xue N. 2008. Labeling Chinese Predicates with Semantic Roles. Computational Linguistics, 34(2):225 255.

Zhuang L., F. Jing, and X. Zhu. 2006. Movie Review Mining and Summarization. In *Proceedings of CIKM 06*, 43 50.