Aspect-Object Alignment Using Integer Linear Programming

Yanyan Zhao[1], Bing Qin[2], and Ting Liu[2]

[1] Department of Media Technology and Art, Harbin Institute of Technology
[2] Department of Computer Science and Technology, Harbin Institute of Technology
{yyzhao,bqin,tliu}@ir.hit.edu.cn

Abstract. Target extraction is an important task in opinion mining, in which a complete target consists of an aspect and its corresponding object. However, previous work always simply considers the aspect as the target and ignores an important element "object." Thus the incomplete target is of limited use for practical applications. This paper proposes a novel and important sentiment analysis task: aspect-object alignment, which aims to obtain the correct corresponding object for each aspect, to solve the "object ignoring" problem. We design a two-step framework for this task. We first provide an aspect-object alignment classifier that incorporates three sets of features. However, the objects assigned to aspects in a sentence often contradict each other. To solve this problem, we impose two kinds of constraints: intra-sentence constraints and inter-sentence constraints, which are encoded as linear formulations and use Integer Linear Programming (ILP) as an inference procedure to obtain a final global decision in the second step. The experiments on the corpora of camera domain show the effectiveness of the framework.

Keywords: Opinion Mining, Aspect-Object Alignment, Integer Linear Programming.

1 Introduction

Opinion mining and sentiment analysis entail a number of interesting and challenging tasks, such as sentiment classification, sentiment extraction and sentiment summarization and so on. A fundamental one is target extraction, which aims to recognize the main topic that has been commented on in a review. Generally speaking, the target is composed of object and aspect. For example, in the review "I bought $[Canon\ 600D]^o$ yesterday, its $[photos]^a$ are amazing," the complete target is $\langle photos, Canon\ 600D\rangle$. However, previous work always considers only the aspect as target [5,16], in which the target is not complete. For instance, in the aforementioned review, the aspect "*photos*" in the second sentence is always directly tagged as the target. Apparently, this target is incomplete because it ignores a very important element, which is the object "*Canon 600D*" that the aspect "*photos*" belongs to. We call this problem as "object ignoring". The incomplete target is of limited use for practical applications if we do not recognize the object.

C. Zong et al. (Eds.): NLPCC 2014, CCIS 496, pp. 193–204, 2014.

In this paper, we define the target consist of two parts, namely, the aspect and its corresponding object, such as $\langle photos, Canon\ 600D\rangle$. Correspondingly, the task of target extraction can be composed of two main subtasks:

– Aspect/object extraction, which aims to extract the aspects/objects in sentiment sentences.
– Find the correct, corresponding object for each aspect in the review. In this paper, we call this task as ***Aspect-Object Alignment***.

Currently, researchers mainly focus on the first task and ignore the second task even though this task is more important for practical applications. The target distribution statistics show that only 10% aspects can explore its objects in the same sentence, which is very low. This result illustrates that just a few aspects and their corresponding objects are co-occurring in the same sentence, such as them in the sentence "the $[appearance]^a$ of $[GF3]^o$ is very beautiful." However, most of the aspects and their corresponding objects do not co-occur, such as them in "the $[appearance]^a$ is very beautiful." Therefore, we should explore the objects for most aspects in other sentences and choose the right one from several object candidates. All the above can indicate that aspect-object alignment is a new and necessary task. As such, this paper mainly focuses on this task. We hypothesize that all aspects and objects appearing in each sentiment sentence have been manually annotated.

We can simply regard the aspect-object alignment task as a binary classification, in which each decision is made in a pair-wise manner made of aspect and object, independently of others. Moreover, we propose three sets of features for the aspect-object classifier, namely, the basic, relational, and special target features. However, there is one major drawback with this method. That is, it makes each decision independently of previous ones in a greedy way. Clearly, the determination of the relation between the aspect and object should be conditioned on how well it works as a whole.

We tackle this issue by recasting the task of aspect-object alignment as an optimization problem, namely, an Integer Linear Programming (ILP) problem. ILP can perform global inference based on the output of the classifier, which makes it highly suitable to address the aforementioned problem. ILP-based models have been developed for many tasks that range from semantic role labeling [11] to multi-document summarization [10], and opinion mining [1]. In this paper, we firstly use ILP to search for a global assignment based on decisions obtained through the binary aspect-object classifier alone. Second, we provide the joint formulation, in which constraints are added to ensure that the ultimate results are mutually consistent.

We construct two kinds of constraints.

Intra-Sentence Constraints: They describe the constraints between objects and aspects or between two aspects, where the objects and aspects appear *in a same sentiment sentence*. For example, if a normal sentence (non-comparative sentence) contains two aspects but no object, then the two aspects share the same object.

Inter-Sentence Constraints: They describe the constraints between objects and aspects or between two aspects, where the objects and aspects appear *in different sentiment sentences*. This kind of constraint always uses the idea of sentiment consistency, which is inspired by the work of Ding et al. [3].

We evaluate our proposed framework on a review corpus of digital camera domain. The experimental results show that the two kinds of constraints achieve significant performances, which are higher than that of the base classifier. The joint model particularly achieves accuracy improvements of more than 5% over the cascading rule-based baseline and nearly 2% over the aspect-object alignment classifier.

2 Related Work

Target extraction is an important task in sentiment analysis, which has recently attracted much attention. Many efficient approaches have been developed for this task, which can be divided into three kinds of methods, namely, rule-based [5,16,12], supervised [14,6], and topic model-based [9,7,8] methods. However, most researchers consider the aspects alone as the targets; thus, the mentioned algorithms are all designed for aspect extraction. The "object" element in the target is ignored, and few researchers are studying on the aspect-object alignment task.

The aspect-object alignment task is similar to the entity assignment task [3], which assigns objects to each sentence in a review. Many rules have been proposed to simply solve this task, which mainly focus on processing the comparative sentences by using the idea of sentiment consistency. The major problem is the sequential application of the rules in this method, which sometimes causes conflicts. Thus, this method cannot effectively achieve an optimal result. This task is actually different from our proposed study, which aims to find an object for each aspect. However, we can simply modify this method as a baseline and use it as a comparative method for our approach.

Aspect-object alignment task is also similar to the task of coreference resolution [13,4] to a particular extent, which has been considerably studied previously. In this task, aspect can be treated as anaphor, and object can be treated as entity. Recently, several researchers studied coreference resolution for the review texts [2]. We can observe that the most significant approach is based on supervised learning, in which a pair-wise function is used to predict a coreferent pair of noun phrases. These methods are an inspiration for us to learn several useful features, such as distance features, for the aspect-object alignment task.

3 Method

This section introduces the proposed overall framework. The framework consists of two main steps, which are learning an aspect-object alignment classifier and using an ILP inference. Generally speaking, an aspect-object alignment classifier is learned to estimate the probability for each pair of aspect and object. We use an ILP inference procedure to achieve an optimal global result by considering specific special constraints.

3.1 Aspect-Object Alignment Classifier

We can generate the Cartesian product of all the aspects and objects in each review into a pair-wise vector $\langle a, o\rangle$. The aspect-object alignment task can be converted to classify each $\langle a, o\rangle$ into true or false. We use a maximum entropy model for the aspect-object alignment classifier.

The features are generated from three kinds of relative sentences, which are shown as follows. Figure 1 shows the examples.

Review1:
s1: [Canon S100]^{o} is really good for general use.
s2: I bought one last year.
s3: The [screen]^{a} is really clear.

Review2:
s1: The [screen]^{a} is really clear.
s2: I strongly recommend [Canon S100]^{o} !

Fig. 1. Example of three kinds of sentences

– Present sentence: this kind of sentence refers to the sentence containing a, such as s_3 of review1 and s_1 of review2 in Figure 1.
– Previous sentence: this kind of sentence should satisfy two conditions: (1) it contains objects and; (2) it is the nearest previous sentence to the present sentence. For example, s_1 of review1 is the previous sentence of the present sentence s_3. We explore the features generated from this kind of previous sentence, because in particular cases, the corresponding object for a in the present sentence can be acquired from the previous sentence.
– Nearest sentence: this kind of sentence should also satisfy two conditions: (1) it contains objects, and (2) it is nearest to the present sentence. Note that, the nearest sentence is different from the previous sentence, because sometimes it can be found in the sentences after the present sentence. We consider the features generated from the nearest sentence because in particular cases, the corresponding object of a in the present sentence can be acquired from the nearest sentence.

Based on the three kinds of sentences, we propose three categories of features, which are basic, relational and special target features, as follows:

Basic Features: We design several basic features from the present, previous and nearest sentence respectively.

– **Sentence Type Feature:** Sentiment sentences can be divided into three types based on the objects or aspects it contains. The first one contains

objects but no aspect. For example, "$[CanonS100]^o$ is really good for general use." The second one contains objects and aspects at the same time. For example, "The $[screen]^a$ of $[CanonS100]^o$ is really clear." The third one contains aspects but no object. For example, "The $[screen]^a$ is really clear." Their possible encoded values are 01, 02 and 03 respectively. This feature is used to describe the present/previous/nearest sentence respectively. We design this feature because the methods of aspect-object alignment vary for different types of sentences.
– **Comparative Sentence Feature:** Sentiment sentences can be divided into normal and comparative sentences. If the present/previous/nearest sentence is a comparative sentence, the value is true; otherwise, is false. We design this feature because the method of aspect-object alignment for normal sentences is different from the method for comparative sentences.
– **Object Feature:** This feature refers to the object that appears in the present/previous/nearest sentence.

Relational Features: We also design several relational features. One is the distance feature, which is inspired by the coreference resolution task [2]. The other one is object consistency feature, which is inspired by Ding et al. [3].

– **Distance between present and previous sentence**: The possible values are 0, 1, 2, 3 and so forth, which captures the sentence numbers between the present and its previous sentence.
– **Distance between present and nearest sentence**: The possible values are 0, 1, 2, 3 and so forth, which captures the sentence numbers between the present and its nearest sentence.
– **Consistency between the object in previous sentence and the candidate object in $\langle a, o\rangle$**: If the object in the previous sentence is the same as the one in $\langle a, o\rangle$, the value is true; otherwise, the value is false.
– **Consistency between the object in nearest sentence and the candidate object in $\langle a, o\rangle$**: If the object in the nearest sentence is the same as the one in $\langle a, o\rangle$, the value is true; otherwise, the value is false.

Special Target Features

– **First appearing object in the review:** This feature refers to the object that appears for the first time in the review.
– **Most frequent object in the review:** This feature refers to the object that appears most frequently in the review.

Based on the above features, we can build an aspect-object alignment classifier to judge each $\langle a, o\rangle$ candidate.

3.2 ILP Inference

In an ideal setting that has a perfect aspect-object alignment classifier, each aspect can obtain the correct object according to the classifier's prediction.

In reality, labels (objects) assigned to aspects in a sentence often contradict each other. For example, each aspect has only one object. These complicated features are difficult to use in the classifier. Therefore, we encode the constraints as linear formulations to resolve the conflicts. We also use ILP as an inference procedure to make a final decision that is consistent with the constraints.

We formally define the aspect set as $A = \{a_1, a_2, ..., a_n\}$, and the object set as $O = \{o_1, o_2, ..., o_m\}$ for each review. We assume that the resulting object set for the aspects in A is $S = \{s_1, s_2, ..., s_n\}$, and $s_i \in O$, which is the resulting object for a_i. Thus, this task can be formulated such that the aspect-object assignment classifier attempts to assign object from O for each a_i in set A. If we assume that the classifier returns a probability value, $p(a_i, s_i)$, which corresponds to the likelihood of assigning label s_i for aspect a_i, then the inference task in a given review can be depicted by maximizing the overall score of the aspects as follows. Moreover, $\hat{S}$ is the optimal result among all the potential vectors.

$$\hat{S} = argmax \sum_{i=1}^{n} p(a_i, s_i) \tag{1}$$

In this equation, the probability $p(a_i, s_i)$ can be obtained through the aforementioned aspect-object alignment classifier in Section 3.1. If s_i is the j^{th} object in the object set O, then $p(a_i, s_i)$ can be denoted as p_{ij}, which represents the probability of the pair of the i^{th} aspect in A and j^{th} object in O.

We can introduce a set of binary indicator variables $z_{ij} \in \{0, 1\}$ for each p_{ij} to reformulate the original $\hat{S}$ function into a linear function, to acquire the optimal $\hat{S}$. If the resulting corresponding object for a_i is actually o_j, then the value $z_{ij} = 1$; otherwise, $z_{ij} = 0$. Thus, the task of finding an optimal $\hat{S}$ can be converted to find an optimal vector Z that can maximize the objective function. Here, Z is the set of z_{ij}, and $i \in \{1, 2, ..., n\}$, $j \in \{1, 2, ..., m\}$. The equation (2) can be written as an ILP objective function as follows:

$$\hat{Z} = argmax \sum_{i=1}^{n} \sum_{j=1}^{m} p_{ij} z_{ij} \tag{2}$$

Subject to

$$\sum_{j=1}^{m} z_{ij} = 1, \qquad \forall z_{ij} \in Z \tag{3}$$

Note that although this constraint comes from the variable transformation, it has a real meaning, which denotes that each aspect can take only one object. This constraint can be considered as **constraint 1**.

Next, we impose several constraints on equation (3) to acquire the optimum solution. The designed constraints can be divided into two categories, which are **intra-sentence constraints** and **inter-sentence constraints** defined in Section 1.

Intra-sentence Constraints. Constraint 2 to Constraint 4 are the representatives.

Constraint 2: If the aspect a_p and a_q are in the same sentence s, and no object appears in the sentence, then a_p has the same object as a_q has. For example, in the sentence "It has a good $[resolution]^a$ and a good $[LCD\ screen]^a$," "*resolution*" and "*LCD screen*" satisfy this constraint; thus, they have the same object. This constraint can be represented using the following equation.

$$\forall j \in \{1, ..., m\}: \qquad z_{pj} = z_{qj}, \tag{4}$$

where two aspects a_p and a_q, but no object appear in s.

Constraint 3: This constraint is available only when the given sentence s can satisfy three conditions: (1) s is a comparative sentence; (2) s contains two objects o_k and o_t, which appear on both sides of the comparative word; and (3) only one aspect a_p appears in this sentence. Then the corresponding object for a_p is one of the two appearing objects in the given sentence. For example, in the sentence "the $[shutter\ sound]^a$ of $[Canon\ 5D3]^o$ is better than $[5D2]^o$'s," the aspect "*shutter sound*" definitely belongs to one of the objects appearing in this sentence. The object in this sentence is "*Canon 5D3*." This constraint can be described by the following formulation:

$$z_{pk} + z_{pt} = 1, \tag{5}$$

where only one aspect a_p and two objects o_k, o_t appear in the comparative sentence s; and meanwhile o_k, o_t appear on both sides of the comparative word.

Constraint 4: This constraint is available only when the given sentence s can satisfy two conditions: (1) s contains only one aspect a_p and one object o_k; and (2) this sentence is a normal sentence. Then the corresponding object for a_p is the object o_k. For example, in the sentence "the $[shutter\ sound]^a$ of $[Canon\ 5D3]^o$ is good," "*Canon 5D3*" is the right object for the aspect "*shutter sound*." This constraint can be described by the following formulation:

$$z_{pk} = 1, \tag{6}$$

where only one aspect a_p and one object o_k appear in the normal sentence s.

Inter-Sentence Constraints. Many previous research illustrate that adjacent sentences in a review have particular sentiment relationships [15,3]. Thus, the sentiment orientations for two adjacent sentences are always the same or totally different (because of the usage of transitional words, such as "but"), which is named as "sentiment consistency." This idea is also very useful for the aspect-object alignment task. We design two constrains based on this idea as follows.

Constraint 5: This constraint is available only when the given sentence s_g can satisfy three conditions: (1) s_g only contains an aspect a_p, but no object; and (2) the previous sentence s_p is a normal sentence, which contains an aspect a' and an object o_k; and (3) the sentiment orientation of s_g is the same as that of s_p. Then the corresponding object for a_p is the object o_k in the previous sentence s_p. For example, in the sentence "The best feature about $[Canon\ S110]^o$ is its $[size]^a$. The $[picture\ quality]^a$ is good, too. ", the second sentence satisfies this constraint, and the corresponding object for the aspect "*picture quality*" is the object "*Canon S110*" in the previous sentence. This constraint can be described by the following formulation:

$$z_{pk} = 1, \tag{7}$$

where only one aspect a_p but no object appear in the given sentence s_g; and only one aspect a' and one object o_k appear in the normal previous sentence s_p; meanwhile, $polarity(s_g) = polarity(s_p)$.

Constraint 6: This constraint is available only when the given sentence can satisfy two conditions: (1) the given sentence s_g contains an aspect a_p, but no object; and (2) the previous sentence s_p is a comparative sentence, and contains two objects o_k and o_t, which appear on both sides of the comparative word and o_k is in front of o_t. If the given sentence shows the same sentiment orientation as the previous sentence, then the corresponding object for a_p is the object o_k in the previous sentence. However, if the given sentence shows a different sentiment orientation with the previous sentence, the object for a_p is o_t.

For example, in the review "the $[shutter\ sound]^a$ of $[Canon\ 5D3]^o$ is better than $[5D2]^o$'s. The $[picture\ quality]^a$ is good, too.", the object for "*picture quality*" in the second sentence is "*Canon 5D3*" in the previous sentence, but not "*5D2*." This constraint can be described by the following formulation:

$$z_{pk} = 1, \qquad if\ polarity(s_g) = polarity(s_p) \tag{8}$$

$$z_{pt} = 1, \qquad if\ polarity(s_g) = -polarity(s_p), \tag{9}$$

where only one aspect a_p but no object appear in the given sentence s_g; and two objects o_k and o_t appear in the previous comparative sentence s_p and on both sides of the comparative word.

The intra-sentence and inter-sentence relationships discussed in the previous sections can be encoded as constraints in an ILP inference process. By formulating the problem this way, we can use the aspect-object alignment classifier and also jointly use an ILP model and many useful constraints to generate the optimal results.

4 Experimental Setup

Corpus. We manually collected on-line customer reviews of digital camera as a case study for the aspect-object alignment task. The corpus is from two famous Chinese forum sites, namely, http://ww.xitek.com/ and http://www.fengniao.com/. The statistics of the corpus are illustrated in Table 1.

Table 1. Statistics for the corpus

No.	Types	Digital camera
1	# of reviews	200
2	# of sentences	8,042
3	# of aspects	2,017
4	Average # of object for each review	2.82
5	# of pairwise $\langle a, o \rangle$	9,161

The raw corpus has 200 documents, in which 8,042 sentences are annotated. We summarized some statistics, which shows 2,017 aspects in the corpus and an average of 2.82 objects for each review. According to the Cartesian product of all aspects and objects, each review in the corpus has a total of 9,161 aspect-object pairs, which are also the input of the binary aspect-object alignment classifier.

Baselines. We compare our system with two baselines. **Baseline1** is a cascading rule-based approach, which is similar to the method of Ding et al. [3]. This approach combines several useful rules including the idea of sentiment consistency, but several rules are conflicting with one another during processing. **Baseline2** is the aspect-object alignment classifier without the ILP inference.

Training and Evaluation. We experiment with ME algorithm and tune the related parameters for the aspect-object alignment classifier by using the 10-fold cross-validation on the corpus described in the previous section. Because our goal of the aspect-object alignment task is to find a certain object for each aspect, the value *Accuracy* is suitable for this task.

5 Results and Discussion

5.1 Results of Our Method and Two Baselines

Table 2 shows the performances of our ILP inference method and two baselines, the cascading rule-based approach, and aspect-object alignment classifier.

Table 2. Comparative results of our method and two baselines

Method	*Accuracy*(%)
Baseline1: cascading rule-based	78.04
Baseline2: aspect-object alignment classifier	81.80
Our ILP inference method	**83.69**

Table 2 shows that the performance of the cascading rule-base approach is not very ideal. The reason is attributed to its nature as a rule based method, where the rules are used sequentially and sometimes have conflicts in them. Baseline2, which is the aspect-object alignment classifier, achieves an accuracy of 81.80%, which significantly (χ test with $p < 0.0001$) outperforms the cascading rule-based approach. The features proposed in this method, which can globally describe this task, provide a major difference.

In addition, our method, which combines an aspect-object alignment classifier and an ILP inference processing, performs best and significantly (χ test with $p < 0.03$) outperforms the two baselines. We can obviously observe that our method with ILP inference performs better that the method (namely, baseline2) without it. This illustrates that the ILP inference is useful, and the two kinds of constraints designed in this paper are effective.

5.2 Aspect-Object Alignment with ILP Inference

ILP is used to perform global inference based on the classifier's output to resolve conflicts between rules. Six constraints are used in this paper, in which Constraint 1 to Constraint 4 reflect the constraints within a sentiment sentence, whereas Constraint 5 and 6 reflect the constraints between different sentiment sentences. Table 3 lists the performance of the system with different constraints.

Table 3. Results of aspect-object alignment with different ILP constraints

Constraints	ILP constraints	*Accuracy* (%)
Intra-sentence constraints	ILP-c1	81.80
	ILP-c2	81.85
	ILP-c3	81.90
	ILP-c4	82.65
Inter-sentence constraints	ILP-c5	82.45
	ILP-c6	82.05
All constraints	ILP-c1-6	**83.69**

Table 3 shows that:

- Constraint 1 refers that there is only one object for each aspect, thus the result with this constraint is equivalent to the basic aspect-object alignment classifier.
- Constraint 2 to Constraint 4 are intra-sentence constraints, and align the aspect-object pair according to the structural information (such as the relationships between aspects and objects) conveyed from the given sentence. The improvement of Constraint 4 among all the constraints is apparent, which illustrates the effectiveness of this constraint. Constraint 2 and 3 can slightly increase the accuracy compared with the basic classifier without ILP inference. That's because most of the instances that satisfy these two constraints can also obtain the correct results by using the basic classifier.
- Constraint 5 and 6 are inter-sentence constraints, and mainly use the idea of sentiment consistency. Table 3 shows that the improvement of Constraint

5 and Constraint 6 are all apparent. These two constraints primarily focus on sentences that contain aspects but no objects. The statistics show that the proportion of this kind of sentences is about 90%, which is very high. The effectiveness of these two constraints demonstrates that for this kind of sentences, we can seek for the corresponding objects in their neighboring sentiment sentences. Moreover, this can further demonstrate that sentiment consistency can improve the task's performance.

- We combine Constraint 1 to Constraint 6 as the final inference and obtain the best performance of 83.69%. It can further improve the aspect-object alignment task by about 2%, which significantly (χ test with $p < 0.03$) outperforms the aspect-object alignment classifier without ILP inference.

6 Conclusion and Future Work

In this paper, we propose a novel and important sentiment analysis task, ***aspect-object alignment***, which aims to resolve the "object ignoring" problem in target extraction. We propose a two-step framework for this task, including an aspect-object alignment classifier and an ILP inference. The experimental results show that the aspect-object alignment classifier with the ILP inference performs better than the classifier without it and also illustrate that the six constraints proposed in this paper are very useful.

In the future, we will endeavor to enhance the algorithm for each step by seeking for more useful features for the classifier or more useful constraints for the ILP inference to improve the performance of the task.

Acknowledgments. We thank the anonymous reviewers for their helpful comments. This work was supported by National Natural Science Foundation of China (NSFC) via grant 61300113, 61133012 and 61273321, and the Ministry of Education Research of Social Sciences Youth funded projects via grant 12YJCZH304.

References

1. Choi, Y., Cardie, C.: Adapting a polarity lexicon using integer linear programming for domain-specific sentiment classification. In: Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing, EMNLP 2009, vol. 2, pp. 590–598 (2009)
2. Ding, X., Liu, B.: Resolving object and attribute coreference in opinion mining. In: Proceedings of the 23rd International Conference on Computational Linguistics, COLING 2010, pp. 268–276. Association for Computational Linguistics, Stroudsburg (2010)
3. Ding, X., Liu, B., Zhang, L.: Entity discovery and assignment for opinion mining applications. In: Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2009, pp. 1125–1134. ACM, New York (2009)

4. Gärtner, M., Björkelund, A., Thiele, G., Seeker, W., Kuhn, J.: Visualization, search, and error analysis for coreference annotations. In: Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations, pp. 7–12 (June 2014)
5. Hu, M., Liu, B.: Mining and summarizing customer reviews. In: Proceedings of KDD 2004, pp. 168–177 (2004)
6. Jakob, N., Gurevych, I.: 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, pp. 1035–1045. Association for Computational Linguistics, Stroudsburg (2010)
7. Liu, K., Xu, L., Zhao, J.: Extracting opinion targets and opinion words from online reviews with graph co-ranking. In: Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 314–324 (June 2014)
8. Mei, Q., Ling, X., Wondra, M., Su, H., Zhai, C.: Topic sentiment mixture: Modeling facets and opinions in weblogs. In: Proceedings of the 16th International Conference on World Wide Web, WWW 2007, pp. 171–180 (2007)
9. Mukherjee, A., Liu, B.: Aspect extraction through semi-supervised modeling. In: Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 339–348. Association for Computational Linguistics, Jeju Island (2012)
10. Nishikawa, H., Hasegawa, T., Matsuo, Y., Kikui, G.: Opinion summarization with integer linear programming formulation for sentence extraction and ordering. In: Proceedings of the 23rd International Conference on Computational Linguistics: Posters, COLING 2010, pp. 910–918. Association for Computational Linguistics, Stroudsburg (2010)
11. Punyakanok, V., Roth, D., Yih, W.T., Zimak, D.: Semantic role labeling via integer linear programming inference. In: Proceedings of the 20th International Conference on Computational Linguistics, COLING 2004, Association for Computational Linguistics, Stroudsburg (2004)
12. Qiu, G., Liu, B., Bu, J., Chen, C.: Opinion word expansion and target extraction through double propagation. Computational Linguistics 37(1), 9–27 (2011)
13. Soon, W.M., Ng, H.T., Lim, D.C.Y.: A machine learning approach to coreference resolution of noun phrases. Comput. Linguist. 27(4), 521–544 (2001)
14. Yu, J., Zha, Z.J., Wang, M., Chua, T.S.: 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, HLT 2011, vol. 1, pp. 1496–1505. Association for Computational Linguistics, Stroudsburg (2011)
15. Zhao, J., Liu, K., Wang, G.: Adding redundant features for crfs-based sentence sentiment classification. In: Proceedings of the Conference on Empirical Methods in Natural Language Processing, EMNLP 2008, pp. 117–126. Association for Computational Linguistics, Stroudsburg (2008)
16. Zhao, Y., Qin, B., Hu, S., Liu, T.: Generalizing syntactic structures for product attribute candidate extraction. In: Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics, pp. 377–380. Association for Computational Linguistics, Los Angeles (2010)