Implicit Polarity and Implicit Aspect Recognition in Opinion Mining

Huan-Yuan Chen and Hsin-Hsi Chen
Department of Computer Science and Information Engineering
National Taiwan University, Taipei, Taiwan
r04922009@ntu.edu.tw; hhchen@ntu.edu.tw

Abstract

This paper deals with a double-implicit problem in opinion mining and sentiment analysis. We aim at identifying aspects and polarities of opinionated statements not consisting of opinion words and aspect terms. As a case study, opinion words and aspect terms are first extracted from Chinese hotel reviews, and then grouped into positive (negative) clusters and aspect term clusters. We observe that an implicit opinion and its neighbor explicit opinion tend to have the same aspect and polarity. Under the observation, we construct an implicit opinions corpus annotated with aspect class labels and polarity automatically. Aspect and polarity classifiers trained by using this corpus is used to recognize aspect and polarity of implicit opinions.

1 Introduction

Opinions are classified into explicit and implicit ones depending on the subjectivity and objectivity (Liu, 2012; Zhang and Liu, 2014). It is more challenging to detect implicit opinions than explicit ones due to the lack of explicit opinion words in the sentences. Aspects refer to facets of the target entities in opinions. They are also categorized into explicit and implicit ones depending on the occurrences of aspect terms. Recognizing implicit aspects in implicit opinions is much more challenging because both opinion words and aspect terms are absent in opinionated statements.

Implicit opinions often describe the situations at which persons concern in their reviews. (S1) and (S2) are two examples selected from positive and negative rating rows respectively in hotel reviews. They do not mention any explicit opinion words and aspect terms. The situation of "many restaurants nearby" infers the convenience for eating, while the situation of "a lot of ants" infers the dirtiness of a room. The implicit opinion describes not only the situation at which customers feel, but also infers the reason why they have such feelings. Implicit opinions are positive in (S1) and negative in (S2), and the implied aspects are location and cleanness.

(S1) 附近有很多餐廳。(There are many restaurants nearby.)

(S2) 房間裡有很多螞蟻。(There are a lot of ants in the room.)

The implicit opinions may be subjective in some cases. For example, (S1) may be placed in negative rating row in a hotel review. Its implicit interpretation will become "There are many restaurants nearby, and thus the air pollution is severe and the smell of the air is very bad."

People may describe a situation first, and then reveal their attitudes and judgments. (S3) is an example. The first clause (only ten meters to the subway entrance) describes a situation, while the second clause (the location is good) is an explicit opinion. In Chinese review, an explicit opinion can also be specified before a situation description. (S4) is an example. In both cases, the polarity and the aspect of the situation are consistent with those of the explicit opinions.

(S3) 到地鐵出入口僅十米，地段好。(Only ten meters to the subway entrance, good location.)

(S4) 地點不錯，可步行至周圍三個捷運站。(Location is good, within walking distance of three MRTs around.)

This paper aims at extracting implicit opinions and identifying their implicit aspects and polarity. We will extract opinions from Chinese hotel reviews, then transfer polarity and aspect from explicit expressions to the corresponding implicit opinions, and train aspect and polarity classifiers. We evaluate the performance of polarity and aspect recognition on implicit opinions.

Almost all previous approaches identify implicit aspects in explicit opinions. They extract opinion words from opinionated sentences, regard them as

Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, pages 20–25,
Berlin, Germany, August 7-12, 2016.

implicit aspect clues, and find aspects through opinion word-aspect term mapping. The lack of opinion words in implicit opinions results in no indicators in mapping. To the best of our knowledge, this paper is the first one to resolve a double-implicit problem in opinion mining and sentiment analysis.

This paper is organized as follows. Section 2 gives a survey on implicit aspect recognition in opinion mining and sentiment analysis. Section 3 constructs an implicit opinions corpus labelled with aspect classes and polarity automatically. Section 4 presents classifiers for implicit polarity and implicit aspect recognition. Section 5 shows and discusses the experimental results.

2 Related Work

Hu and Liu (2004) present the first feature-based opinion summarization system. They point out explicit and implicit product features, and extract explicit features by using association miner and pruning strategies. The opinionated sentences along with their polarity are listed under individual product features. Popescu and Etzioni (2005) introduce an opinion extraction system OPINE. OPINE extracts explicit product features based on Point-wise Mutual Information. This work does not discuss the implicit feature generation. Liu et al. (2005) present an association mining approach to extract both explicit and implicit features in their opinion observer, but the implicit features discussed occur explicitly in an overt form, e.g., [MB] indicates a product feature <memory>.

Su et al. (2008) define an implicit feature as the product feature which does not occur explicitly, but can be inferred from the surrounding opinion word. They propose a mutual reinforcement approach to cluster product features and opinion words simultaneously, and extract implicit features based on opinion words. In the subsequent work, different methodologies are proposed to identify the association between opinion words and aspect terms (called also product features), thus implicit aspects are inferred from opinion word-aspect term mapping (Bagheri et al., 2013).

Zhen et al. (2011) propose a two-phase co-occurrence association rule mining approach. Yu et al. (2011) generate a review hierarchy based on aspects. Implicit aspect of a review can be determined by the cosine similarity of the review vector and the vector for each aspect node in the review hierarchy. Zeng and Li (2013) regard identification of implicit features as a classification problem, and consider reviews for each clustered opinion-pair as training set. Wang et al. (2013) employ five collocation methods including frequency, PMI, frequency/PMI, t-test and chi-square test to measure the association between opinion words and aspect terms.

Cruz et al. (2014) manually annotate implicit aspects and implicit aspect indicators (IAI) on the customer review datasets in Hu and Liu (2004), and employ Conditional Random Fields to recognize IAI. Poria et al. (2014) identify implicit aspect clues (IACs) in a document. Both approaches establish IAI (IAC) and aspect mapping.

Mukherjee and Liu (2012) propose two statistical models to deal with aspect categorization problem. They use hotel reviews from tripadvisor.com, and point out categorizing aspects is a subjective task. Total 9 major aspects based on commonsense knowledge, including Dining, Staff, Maintenance, Check In, Cleanliness, Comfort, Amenities, Location and Value for Money, are considered. Kim et al. (2013) further analyze general aspects and specific aspects, and discuss how aspect structure is helpful. Zhao et al. (2015) present a fine-grained corpus for sentiment analysis.

Our work is different from the previous ones in two-fold: (1) opinion is implicit, so that no opinion words can be used as clues; and (2) aspect is implicit, so that no aspect terms can be found. The direct opinion word and aspect mapping is not feasible in implicit polarity and implicit aspect recognition. We focus on the construction of an implicit opinions corpus for double-implicit recognition. The aspect categorization is not the major concern.

3 Constructing Implicit Opinions Corpus

This section first defines the implicit opinions, collects a Chinese hotel dataset, identifies opinion and aspect clusters from the dataset, and constructs implicit opinion corpus labelled with aspect class and polarity.

3.1 Definitions of Implicit Opinions

A sentence in a review can be partitioned into several segments separated by punctuation marks. The

following show four possible types of segments based on the occurrences of opinion words and aspect terms, where + and - denote occurrence and non-occurrence. Segments of types (T1) and (T2) contain explicit opinion words, while segments of types (T3) and (T4) contain no opinion words. They appear together with and without aspect terms.

(T1) (+opinion word, +aspect term)
e.g., 地點不錯 (location is good)
(T2) (+opinion word, -aspect term)
e.g., 很便宜 (very cheap)
(T3) (-opinion word, +aspect term)
e.g., 地理位置 (location)
(T4) (-opinion word,-aspect term)
e.g., 到油麻地地鐵站只要兩分鐘 (Just two minutes to Yau Ma Tei MRT Station)

Segments of either type can not only appear individually, but also can be combined with other types of segments to form a sentence. Segments of types (T1) and (T2) are opinionated. Segments of type (T3) are opinionated implicitly when they appear in positive/negative rating row. Segments of type (T4) can be opinionated or non-opinionated. It is interpreted as an opinionated segment clearly when it is placed in rating row individually.

(S5) is a sentence consisting of 5 segments of types T3, T2, T1, T4 and T3, respectively. The 4th segment, i.e., feeling a little like shanty towns, is a double-implicit opinion. Its polarity and aspect (negative and environment) can be inferred from the 3rd segment, i.e., the surrounding environment is really bad.

(S5) [$_{T3}$ 旅館在小巷子裡]，[$_{T2}$ 安全沒有問題]，[$_{T1}$ 但附近環境確實不好]，[$_{T4}$ 有點棚戶區的感覺]，[$_{T3}$ 周圍沒有飯店]。([$_{T3}$ hotel in the alley]，[$_{T2}$ security is no problem]，[$_{T1}$ but the surrounding environment is really bad]，[$_{T4}$ feeling a little like shanty towns]，[$_{T3}$ no hotels around])

In this paper, we deal with opinionated segments of type (T4). On the one hand, we extract pairs of segments of types T1-T4 or T4-T1 from a Chinese hotel review dataset. The segments of type T4 will be annotated with opinion words and aspect terms extracted from their paired segments of type T1. The segments of type T4 along with their annotations form a training corpus. On the other hand, the test segments of types (T4) will be labelled with polarity and aspect by polarity and aspect classifiers.

At first glance, we do not need to perform the classification task on T4 segments since we can directly use polarity and aspect of T1 segments. The scenario is just for test purpose because we do not have large-scale manually-labelled data. In the latter experiments, we will also consider the cases of T4 segments existing individually in rating rows. That will reflect the real situations.

3.2 Extraction of Implicit Opinions

Opinion words and aspect terms are the indicators to define the four types (T1)-(T4). As a case study, we collect a Chinese hotel review dataset from booking.com. It consists of 144,158 positive reviews and 113,844 negative reviews about 20,973 hotels from 49 international cities. Here only Chinese reviews are kept. We use Stanford NLP tools to segment, POS tag, and parse all the reviews.

At first, we construct an opinion dictionary from this dataset. Words of POS tags VA, VV, AD, and JJ are candidates of opinion words. We adopt Chi-square test and point-wise mutual information to filter out less confident words from the candidate set, respectively. We examine the union of the remaining words manually and construct an opinion dictionary consisting of 374 positive and 408 negative opinion words.

Then, we construct an aspect dictionary based on opinion words. A word meeting the following four conditions is regarded as an aspect term candidate: (1) its POS is NN, (2) it occurs at least 100 times, (3) it is accompanied with an opinion word within the same segment, and (4) their dependency is nsubj. We examine 183 proposed candidates manually and construct an aspect dictionary consisting of 153 aspect terms.

In an extreme case, a review does not contain any opinion words and aspect terms. It may be a single segment or multiple segments of type T4. Reviews are listed under positive and negative rating rows, so we know their polarity, but not aspect. Table 1 shows the statistics of such kinds of reviews in the hotel dataset. Interestingly, 2.07% of positive reviews are pure T4, and 7.29% of negative reviews are pure T4. That demonstrates double-implicit is a practical issue and customers tend to express negative opinions implicitly. The pure

	single	multiple	total
# pure T4 (positive reviews)	2,266	717	2,983
# pure T4 (negative reviews)	5,847	2,451	8,298

Table 1: Statistics of pure T4 reviews.

	T1	T2	T3	T4
total	192,353	161,863	257,831	303,357
ratio	21.01%	17.68%	28.17%	33.14%

Table 2: Statistics of segment types.

T4 reviews set consisting of single segments only is called PT4S hereafter.

Table 2 shows the statistics of segments of types T1, T2, T3, and T4. Only 21.01% of segments contain both opinion words and aspect terms, and 33.14% of segments do not contain any opinion words and aspect terms. We further examine the type combinations of two successive segments. There are 103 possible punctuation marks between any two segments, including common ones like “ , ”, “ 。”, “?”, and “!”, and some special ones like “~~~”. To avoid misinterpretation of the special marks, we considers only those segment pairs linked by commas. Moreover, to obtain an automatically labelled dataset, the ambiguous sequence of segments, X-T4-Y, where X and Y of types T1, T2, or T3, are removed. Total 31,136 T4-T1/T1-T4 segment pairs remain. They are used to derive an implicit opinions corpus for learning and testing polarity classifier and aspect classifiers. This data set is called T41 hereafter.

In most of the cases we observed, segment of type T2 or T3 does not pass its aspect or opinion to nearby segments of type T4. (S6) is an example of a triple of segments of type T1-T4-T3, which introduces ambiguity between aspect and opinion assignment. The aspect of segment of type T1, i.e., *the equipment*, competes with that of segment of type T3, *the toilet*. In this case, *the safety deposit box*, which is the undetected aspect of the segment of type T4, and *the toilet* are two sub-aspects of *the equipment*. The latter two clauses are supplementary description of the first clause.

(S6) 設施比較舊，保險箱不好使，馬桶上水時有故障 (The equipment is old, the safety deposit box is hard to use, and the toilet sometimes stucks while refilling.)

This work bases on the postulation – say, an implicit opinion and its neighbor explicit opinion tending to have the same aspect and polarity, to construct a training corpus automatically. We randomly sampled 1% of pairs of segments of type T1-T4 or T4-T1 in a training corpus (see Section 4) to verify whether our assumption holds. In this setup, we discard clauses that contain parsing errors and those are too short to represent both aspects and opinions. The result is promising. On average, 70.46% of the pairs follow the observation. In particular, the pairs keep the property more often (i.e., 74.51%) when the polarity of T1 is negative.

4 Double-Implicit Opinion Analysis

We assign polarity and aspect of a T4-type segment in T41 dataset based on the information from its paired T1-type segment. Negation in the T1-type segment will reverse the polarity. To avoid data sparseness, 153 aspect terms are partitioned into 10 aspect classes based on common sense knowledge, including food, hotel, price, room, internet, staff, services, facilities, neighborhood, and general. The criterion in the selection of the category of aspects is not the major concern in this paper. For example, facilities and services may be merged into the same aspect category. The 31,136 labelled T4-type segments in T41 dataset are divided into training and test sets consisting of 23,352 and 7,784 segments, respectively.

Figure 1 shows the segment length distribution of T41-train, T41-test, T41, and PT4S datasets. The length is measured by number of Chinese words in a segment. X-axis and Y-axis denote length of segments and ratio, respectively. Segments in PT4S dataset are shorter than those in T41 dataset. Segments of 2 and 3 words occupy 48.61%. Table 3 shows the polarity distribution in these datasets. Because T41 dataset is divided into T41-train and T41-test datasets uniformly, their polarity distribution is the same, i.e., positive:

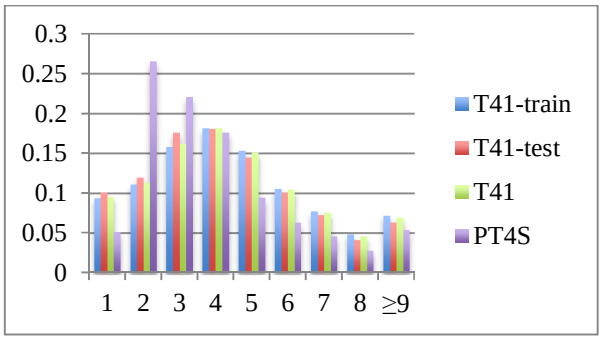

Figure 1: Length distribution in experimental datasets.

	T41	T41-train	T41-test	PT4S
positive	79.64%	79.63%	79.68%	27.93%
negative	20.36%	20.37%	20.32%	72.07%

Table 3: Polarity distribution in experimental datasets.

(%)	BOW linear	W2V linear	BOW RBF	W2V RBF	W2V CNN
T41-test (p)	78.55	73.67	81.54	79.76	**85.04**
PT4S (p)	**77.30**	77.64	72.01	72.22	67.96
MicroAvg	**77.91**	75.69	76.67	75.91	76.32
T41-test (a)	43.25	41.50	46.35	46.13	**55.90**

Table 4: Accuracy of implicit polarity and aspect recognition.

negative=4:1. Comparatively, positive:negative= 1:2.58 in PT4S dataset. The two test sets bias toward different polarities.

We employ T41-train dataset to train binary polarity classifier and 10-way aspect classifiers, and test on T41-test dataset. We also explore T41 dataset to train polarity classifier, and test on PT4S dataset. T41-testing evaluates both implicit polarity and implicit aspect recognition. Note the ground truth is generated automatically. PT4S-testing evaluates implicit polarity only based on the human-annotated ground truth.

We consider bag of words (BOW) and word vectors generated by word2vec (W2V) as features, where word vectors are pre-trained by using the part-of-tagged Chinese sentences extracted from the ClueWeb09 dataset (CMU, 2009; Yu et al., 2012). Moreover, we adopt SVM with linear kernel and SVM with RBF kernel learning algorithms in Scikit-Learn library (Pedregosa et al., 2011), and run cross-validation multiple times on the training set to facilitate a grid search on hyperparameters with F-measure as the metric to optimize.

Besides, we also explore Convolutional Neural Networks (CNN) (Kim, 2014). Table 4 summarizes the accuracy of implicit polarity and implicit aspect recognition, where (p) and (a) after dataset denote polarity and aspect performance of that dataset, respectively. CNN achieves the best implicit polarity and aspect recognition in T41-test dataset. However, its implicit polarity accuracy is decreased to 67.96%. It may be due to overfitting in small amount of training data. Different dropout rates (Srivastava et al., 2014) can be explored. SVM with linear kernel (BOW) gets the best micro average accuracy (77.91%) in implicit polarity recognition.

Figure 2 shows the accuracies of the implicit polarity recognition on segments of different lengths.

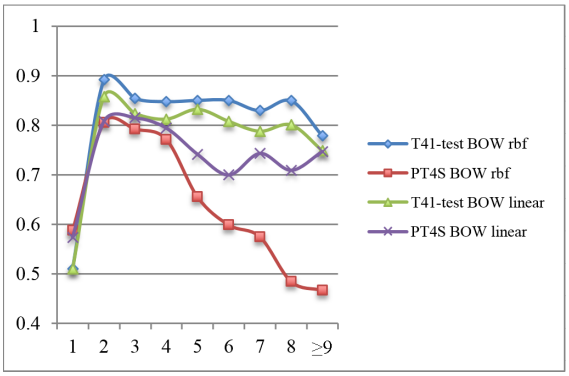

Figure 2: Accuracies of segments of different lengths.

It is challenging to predict the implicit polarity and aspect for segments of very short length. Figure 1 depicts one-word segments occupy 5%-10%. One word segment like "旺角" (Mong Kok) is ambiguous. If we neglect such segments, the micro average accuracy in implicit polarity recognition using SVM with linear kernel (BOW) is increased to 79.94%, and the accuracy in implicit aspect recognition (10-way classification) becomes 46.01%.

5 Conclusion and Future Work

In this paper, we address the double-implicit issue in opinion mining and sentiment analysis, and propose a protocol to derive a labelled corpus for implicit polarity and implicit aspect analysis. SVM with linear kernel (BOW) is robust in implicit polarity recognition. Ten-way classification for implicit aspect recognition still has space to improve.

This work bases on the aspect-and-polarity-transfer postulation to construct a training corpus automatically. We randomly sample T4 segments from T4-T1 or T1-T4 pairs and check them manually. We find that 70.46% of the pairs follow the observation. The experimental setup is reasonable for evaluation with PT4S dataset because it is labelled by users themselves. To derive a more reliable training set, distinguishing if T4 is non-opinionated needs to be investigated further.

Moreover, we neglect the cases T4-X (X-T4), where X is either T2 or T3, in the selection of training set. It is also challenging when either opinion word or aspect term is absent from the cue segment. In this paper, we provide some case studies of these scenarios, but how to utilize the partial information in implicit polarity and implicit aspect recognition is a future work.

Acknowledgments

This research was partially supported by Ministry of Science and Technology, Taiwan, under grant MOST-102-2221-E-002-103-MY3. We thank the anonymous reviewers for their constructive comments to revise this paper.

References

Arjun Mukherjee and Bing Liu. 2012. Aspect Extraction through Semi-Supervised Modeling, *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics*, pages 339–348.

Ayoub Bagheria, Mohamad Saraeeb, and Franciska de Jong. 2013. Care More about Customers: Unsupervised Domain-independent Aspect Detection for Sentiment Analysis of Customer Reviews, *Knowledge-Based Systems*, 52:201–213.

CMU. 2009. ClueWeb09, http://lemurproject.org/clueweb09.php/.

Ivan Cruz, Alexander Gelbukh, and Grigori Sidorov. 2014. Implicit Aspect Indicator Extraction for Aspect-based Opinion Mining, *International Journal of Computational Linguistics and Applications*, 5(2):135-152.

Minqing Hu and Bing Liu. 2004. Mining and Summarizing Customer Reviews, *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 168–177.

Yoon Kim. 2014. Convolutional Neural Networks for Sentence Classification, *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing* (*EMNLP*), pages 1746–1751.

Suin Kim, Jianwen Zhang, Zheng Chen, Alice Oh, and Shixia Liu. 2013. A Hierarchical Aspect-Sentiment Model for Online Reviews, *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence*, pages 526–533.

Bing Liu. 2012. *Sentiment Analysis and Opinion Mining*. Morgan & Claypool Publishers.

Bing Liu, Minqing Hu, and Junsheng Cheng. 2005. Opinion Observer: Analyzing and Comparing Opinions, *Proceedings of the 14th International Conference on World Wide Web*, pages 1024–1025.

Fabian Pedregosa, Gael Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Edouard Duchesnay. 2011. Scikitlearn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.

Ana-Maria Popescu and Oren Etzioni. 2005. Extracting Product Features and Opinions from Reviews, *Proceedings of Conference on Empirical Methods in Natural Language Processing*, pages 3–28.

Soujanya Poria, Erik Cambria, Lun-Wei Ku, Chen Gui and Alexander Gelbukh. 2014. A Rule-Based Approach to Aspect Extraction from Product Reviews, *Proceedings of the Second Workshop on Natural Language Processing for Social Media* (*SocialNLP*), pages 28–37.

Nitish Srivastava, Georey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting, Journal of Machine Learning Research, 15, pages, 1929–1958.

Qi Su, Xinying Xu, Honglei Guo, Zhili Guo, Xian Wu, Xiaoxun Zhang, Bin Swen, and Zhong Su. 2008. Hidden Sentiment Association in Chinese Web Opinion Mining, *Proceedings of International Conference on World Wide Web*, pages 959–968.

Wei Wang, Hua Xu, and Wei Wan. 2013. Implicit Feature Identification via Hybrid Association Rule Mining, *Expert Systems with Applications*, 40(9): 3518–3531.

Jianxing Yu, Zheng-Jun Zha, Meng Wang, Kai Wang, and Tat-Seng Chua. 2011. Domain-Assisted Product Aspect Hierarchy Generation: Towards Hierarchical Organization of Unstructured Consumer Reviews, *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 140–150.

Lingwei Zeng and Fang Li. 2013. A Classification-Based Approach for Implicit Feature Identification, *Proceedings of the China National Conference on Computational Linguistics*, LNAI 8202, pages 190–202.

Lei Zhang and Bing Liu. 2014. Aspect and Entity Extraction for Opinion Mining, *Data Mining and Knowledge Discovery for Big Data*, Studies in Big Data, 1, pages 1-40.

Yanyan Zhao, Bing Qin, and Ting Liu. 2015. Creating a Fine-Grained Corpus for Chinese Sentiment Analysis, *IEEE Intelligent Systems*, 30(1):36–43.

Chi-Hsin Yu, Yi-jie Tang, and Hsin-Hsi Chen. 2012. Development of a Web-Scale Chinese Word N-gram Corpus with Parts of Speech Information, *Proceedings of 8th International Conference on Language Resources and Evaluation*, pages 320–324.

Hai Zhen, Kuiyu Chang, and Jung-jae Kim. 2011. Implicit Feature Identification via Co-occurrence Association Rule Mining, *Proceedings of 12th International Conference on Computational Linguistics and Intelligent Text Processing*, pages 393–404.