

Aspect Based Sentiment Analysis: Category Detection and Sentiment Classification for Hindi

Md Shad Akhtar, Asif Ekbal, and Pushpak Bhattacharyya

Department of Computer Science & Engineering
Indian Institute of Technology Patna
India-801103
{shad.psc15,asif,pb}@iitp.ac.in

Abstract. E-commerce markets in developing countries (e.g. India) have witnessed a tremendous amount of user's interest recently. Product reviews are now being generated daily in huge amount. Classifying the sentiment expressed in a user generated text/review into certain categories of interest, for example, *positive* or *negative* is famously known as sentiment analysis. Whereas aspect based sentiment analysis (ABSA) deals with the sentiment classification of a review towards some aspects or attributes or features. In this paper we assess the challenges and provide a benchmark setup for aspect category detection and sentiment classification for Hindi. Aspect category can be seen as the generalization of various aspects that are discussed in a review. As far as our knowledge is concerned, this is the very first attempt for such kind of task involving any Indian language. The key contributions of the present work are two-fold, viz. providing a benchmark platform by creating annotated dataset for aspect category detection and sentiment classification, and developing supervised approaches for these two tasks that can be treated as a baseline model for further research.

Keywords: Aspect Category Detection, Sentiment Analysis, Hindi

1 Introduction

With the globalization of internet over the past decade or so, usage of e-commerce as well as social media has increased enormously. Users do express their opinions regarding a product and/or service online. Organizations and other users treat these feedbacks and opinions as a goodness measure for the product or service. The amount of contents generated daily poses several practical challenges to maintain and analyze these effectively. Some of the challenges are due to the informal nature of texts, code-mixing (mixing of several language contents) behaviors and the non-availability of many basic resources and/or tools for the processing of these kinds of texts. Thus, it has been a matter of interest to the researchers worldwide to develop robust techniques and tools in order to effectively and accurately analyze the user generated contents. One such task is famously known as sentiment analysis [1] that deals with finding an un-biased

opinion of review or text written in social media platforms. It tends to classify a piece of user written text by predicting its polarity as either positive or negative. Finding the polarity of a user review with respect to some features or aspects is known as aspect based sentiment analysis (ABSA), which is gaining interest to the community because of its practical relevance. In 2014, a SemEval shared task [2] was contributed to address this problem in two domains namely, *restaurant* & *laptop*. It includes four subtasks:

1. Aspect Term Extraction (ATE)
2. Aspect Term Sentiment (ATS)
3. Aspect Category Detection (ACD)
4. Aspect Category Sentiment (ACS)

The first subtask i.e. aspect term extraction, can be thought of as a sequence labeling problem, where for given sequence of tokens, one has to mark the boundary of an aspect term properly. The second problem was a classification problem, where the sentiment expressed towards an aspect has to be classified as positive, negative, neutral and conflict. The problem of aspect category detection (the third task) deals with the classification of an aspect term into one of the predefined categories. The problem related to the fourth task was to classify the sentiment expressed in a review with respect to the aspect category. The third and the fourth tasks in SemEval considered the reviews of only the restaurant domain, and five aspect categories (i.e. food, price, service, ambiance and misc) were defined. Table 1 shows one example review, each for English and Hindi. The English review contains one aspect term i.e. *'bread'* which belongs to the aspect category *'Food'*. Polarities towards both the aspect term and aspect category are *'positive'*. Similarly, Hindi review contains one aspect term i.e. *हाउसिंग* (*haaUsiNg*) and its sentiment is *'neutral'*. However, it belongs to two different aspect categories i.e. *'Design'* & *'Misc'*, and the sentiments towards these are *'neutral'* and *'negative'*, respectively.

Such a fine-grained analysis provides greater insight to the sentiments expressed in the written reviews. In recent times, there have been a growing trends for sentiment analysis at the more fine-grained level, i.e. for aspect based sentiment analysis (ABSA). Few of the interesting systems that have emerged are [3–7]. However, all these research are related to some specific languages, predominantly for English.

Sentiment analysis in Indian (especially Hindi) languages are still largely unexplored due to the non-availability of various resources and tools such as annotated corpora, lexicons, Part-of-Speech (PoS) tagger etc. Existing works [8–16] involving Indian languages mainly discuss the problems of sentiment analysis at the coarse-grained level with the aims of classifying sentiments either at the sentence or document level. Existing works have limited scope, mainly because of the lack of good quality resources and/or tools. For example, Balikwal et.

¹ Transliterated and translated forms are provided only for representation purpose. We did not include them for model construction.

Subtasks	Review Text	
	<i>"The bread is top notch as well".</i>	
ATE	<i>bread</i>	
ATS	<i>positive</i>	
ACD	<i>Food</i>	
ACS	<i>positive</i>	
Subtasks	Review Text	
	Devanagri	``इसका हाउसिंग स्टेनलेस स्टील से निर्मित है इसलिए बहुत भारी है।".
	Transliterated	``Isakaa haaUsiNg sTenales sTeel se nirmit hai IsaliE bahut bhaaree hai.".
	Translated	``Its housing is made up of stainless steel that why it is very heavy.".
ATE	हाउसिंग (<i>haaUsiNg</i>)	
ATS	<i>neutral</i>	
ACD	<i>Design, Misc</i>	
ACS	<i>neutral, negative</i>	

Table 1: Examples of various subtasks of aspect based Sentiment Analysis. ATE: Aspect term extraction; ATS: Aspect term sentiment; ACD: Aspect category detection; ACS: Aspect category sentiment. ¹

al. [11] used Google translator to generate the dataset, which clearly does not guarantee good quality because of the translation errors encountered. On the other hand, the works reported in [9, 10, 8] used the datasets that are limited in size (few 100s reviews). Aspect based sentiment analysis (ABSA) in Indian languages have not been attempted at large-scale so far. Hence, the problem is still an open challenge, mainly, because of the non-availability of any benchmark setup that could provide a high-quality dataset, baseline model as well as the proper evaluation metrics. In recent time, a framework for aspect based sentiment analysis for Hindi has been proposed in [17] that provides annotated dataset for aspect term extraction and sentiment classification with respect to the aspect term.

It provides 5,417 user reviews collected from 12 domains. In this work, our focus is to provide a benchmark framework for aspect category detection and its polarity classification. We create a dataset annotated with aspect categories and their polarities. In order to show the effective usage of the generated dataset we develop models based on supervised approaches for solving two problems, *viz.* aspect category detection and sentiment classification.

The rest of the paper is structured as follows. Section 2 discusses the various aspects of the datasets. Methodologies of aspect category detection and its sentiment classification are described in Section 3. Experimental results along with necessary analysis are presented in Section 6. Finally, in Section 5 we present the concluding remarks.

2 Benchmark setup for ABSA in Hindi

For ABSA there is no available dataset for the Indian languages, in general, and Hindi, in particular. We create our own dataset for aspect category detection and sentiment classification by collecting user generated web reviews, and annotating these using a pre-defined set of categories. In subsequent subsections we describe these steps in details.

2.1 Data Collection

We crawl various online sources² and collect 5,417 user generated reviews, which belong to 12 different domains, namely i) Laptops, ii) Mobiles, iii) Tablets, iv) Cameras, v) Headphones, vi) Home appliances, vii) Speakers, viii) Televisions, ix) Smart watches, x) Mobile apps, xi) Travels and xii) Movies. Details of these dataset statistics are presented in Section 2.3.

2.2 Data Annotation

We define and compile a list of aspect categories for different domains as listed in Table 2. All electronics products or domains (*except* Mobile apps, Travels and Movies) share six common categories among themselves e.g. Design of the product, Software, Hardware, Ease of use or accessibility, Price of the product and Miscellaneous.

We follow similar scheme in line with SemEval shared task for annotating the dataset. We identify various aspect categories of each review along with its associated sentiment and save them into a XML format. Table 3 lists xml structure of two such instances from the dataset. The upper half of the table contains two example reviews in Devanagari script, its Roman transliterated as well as English translated forms. Both the reviews have one aspect category associated with them and whose polarities are *neutral* and *negative*, respectively.

² List of few sources...

<http://www.jagran.com>
<http://www.gizbot.com>
<http://www.patrika.com>
<http://www.hi.themobileindian.com>
<http://www.mobilehindi.com>
<http://navbharattimes.indiatimes.com>
<http://hindi.starlive24.in/>
<http://www.amarujala.com>
<http://techjankari.blogspot.in>
<http://www.digit.in>
<http://khabar.ndtv.com/topic>
<http://www.hindi.mymobile.co.in/>
<http://www.bhaskar.com>

Domains	Aspect Categories
Electronics (Laptops, Mobiles, Tablets, Cameras, Speakers, Smart watches, Headphones, Home appliances & Televisions)	Design, Software, Hardware, Ease of use, Price, Misc.
Mobile apps	GUI, Ease of use, Price, Misc.
Travels	Scenery, Place, Reachability, Misc.
Movies	Story, Performance (Action/Direction etc.), Music, Misc.

Table 2: Aspect categories that correspond to different domains.

The $\langle sentences \rangle$ node represents root node of the xml that contains every sentence of the review as its children i.e. $\langle sentence \rangle$. To uniquely identify each $\langle sentence \rangle$, an $\langle id \rangle$ is associated with it as an attribute. Each $\langle sentence \rangle$ node has three children, namely $\langle text \rangle$, $\langle aspectTerms \rangle$ and $\langle aspectCategories \rangle$. The $\langle text \rangle$ node holds one review sentence, whereas $\langle aspectTerms \rangle$ contains n $\langle aspectTerm \rangle$ nodes as its children if a review sentence has n aspect terms. For the example at hand n equals to 1 and 0 for sentence ids 1 and 2, respectively. Each $\langle aspectTerm \rangle$ node holds four attributes: $\langle term \rangle$, $\langle from \rangle$, $\langle to \rangle$ & $\langle polarity \rangle$. Attribute $\langle term \rangle$ defines aspect term represented by current node while $\langle polarity \rangle$ stores the sentiment towards the $\langle term \rangle$. Position of the aspect term in the review text is determined by attributes $\langle from \rangle$ and $\langle to \rangle$ which store the index of first and last character, respectively in the review text. Similarly, $\langle aspectCategories \rangle$ contains m $\langle aspectCategory \rangle$ nodes if a review belongs to m different categories. Both the review sentences discuss about one category each. The $\langle aspectCategory \rangle$ node has two attributes i.e. $\langle category \rangle$ & $\langle polarity \rangle$ which store the aspect category and its sentiment polarity, respectively.

2.3 Dataset statistics

The dataset contains 5,417 user reviews related to the product or service. There are total of 2,250 *positive*, 635 *negative*, 2,241 *neutral* and 128 *conflict* instances of aspect categories. Overview of the dataset statistics are presented in Table 4.

3 Methodologies for Aspect Category Detection and Sentiment Classification

Aspect category is a high level abstract representation (summarized form) of the aspect terms. In other words, each aspect term must belong to one of the pre-defined categories which represent that aspect term. However, aspect category

Id	Format	Review Text
1.	Devanagari	इसकी स्क्रीन 15.6 इंच की है।
	Transliterated	Isakee skreen 15.6 INch kee hai.
	Translated	It has 15.6 inch screen.
2.	Devanagari	यह बहुत महंगा है।
	Transliterated	yah bahut mahaNgaai hai.
	Translated	It is very costly.
Annotation Structure		
<pre> <sentences> <sentence id= "1" > <text> इसकी स्क्रीन 15.6 इंच की है।< \text> <aspectTerms> <aspectTerm from="5" to="10" term="स्क्रीन" polarity="neutral" /> < \aspectTerms> <aspectCategories> <aspectCategory category="hardware" polarity="neutral" /> < \aspectCategories> < \sentence> <sentence id= "2"> <text> यह बहुत महंगा है।< \text> <aspectCategories> <aspectCategory category="price" polarity="negative" /> < \aspectCategories> < \sentence> <sentence id= "3"> ... < \sentence> < \sentences> </pre>		

Table 3: Dataset annotation structure.

can be implicit as well. A review that does not contain any explicit aspect term can still belong to one of the categories. For e.g., in Table 3, second sentence does not have any aspect term but still it talks about the ‘*price*’ category whose polarity is *negative*. This information is implicitly present in the review because of the occurrence of word महंगा (mahaNgaai=costly). In order to show the efficacy of the resource that we created, we build two separate models for aspect category detection and sentiment classification based on supervised machine learning approaches. We make use of language independent features for both the tasks,

Domains	Polarity	Category													
		HW	SW	Des.	Pri.	Ease	GUI	Place	Rea.	Scs.	Story	Perf	Music	Misc	Total
Electronics (Laptops, Mobiles, Tablets, Cameras, Headphones, HomeApps, Speakers, Smartwatches & Televisions)	Pos	700	160	305	110	70	-	-	-	-	-	-	-	290	1635
	Neg	261	55	69	31	19	-	-	-	-	-	-	-	89	524
	Neu	763	149	137	83	30	-	-	-	-	-	-	-	173	1335
	Conf	73	6	13	4	3	-	-	-	-	-	-	-	21	120
	Total	1797	370	524	228	122	-	-	-	-	-	-	-	573	3614
Mobile Apps	Pos	-	-	-	4	18	14	-	-	-	-	-	-	64	100
	Neg	-	-	-	0	4	5	-	-	-	-	-	-	13	22
	Neu	-	-	-	6	3	8	-	-	-	-	-	-	57	74
	Conf	-	-	-	0	1	0	-	-	-	-	-	-	0	1
	Total	-	-	-	10	26	27	-	-	-	-	-	-	134	197
Travels	Pos	-	-	-	-	-	-	195	7	97	-	-	-	57	356
	Neg	-	-	-	-	-	-	5	9	1	-	-	-	6	21
	Neu	-	-	-	-	-	-	103	19	24	-	-	-	41	187
	Conf	-	-	-	-	-	-	1	0	0	-	-	-	0	1
	Total	-	-	-	-	-	-	304	35	122	-	-	-	104	565
Movies	Pos	-	-	-	-	-	-	-	-	-	6	109	14	30	159
	Neg	-	-	-	-	-	-	-	-	-	11	35	5	17	68
	Neu	-	-	-	-	-	-	-	-	-	17	95	8	525	645
	Conf	-	-	-	-	-	-	-	-	-	1	5	0	0	6
	Total	-	-	-	-	-	-	-	-	-	35	244	27	572	878
Overall	Pos	700	160	305	114	88	14	195	7	97	6	109	14	441	2250
	Neg	261	55	69	31	23	5	5	9	1	11	35	5	125	635
	Neu	763	149	137	89	33	8	103	19	24	17	95	8	796	2241
	Conf	73	6	13	4	4	0	1	0	0	1	5	0	21	128
	Total	1797	370	524	238	148	27	304	35	122	35	244	27	1383	5254

Table 4: Dataset statistics. Pos: positive, Neg: Negative, Neu: Neutral, Conf: Conflict

i.e. we do not use any domain-specific resources or tools for implementing the features.

3.1 Aspect Category Detection

The problem of aspect category detection can be modelled with the multi-label classification framework, where each review belongs to zero (0) or more categories. In general, a multi-label classification problem can be solved using two techniques, such as: i) binary relevance approach and ii) label powerset approach. Binary relevance approach handles the multi-label scenario by first building n distinct models for each n unique label. The prediction of n models are then combined to produce the final prediction. Whereas, label powerset approach treats each label combination as a unique label. It then trains and evaluates the model. An example scenario is depicted in Table 5 for both the approaches. First two rows list 5 text reviews T_i , for $i = 1..5$ and the corresponding class labels. The two-class labels i.e. ‘a’ and ‘b’ can be assigned to any review. For

instance reviews T_1 and T_4 belong to both ‘a’ and ‘b’ classes. In binary relevance approach two separate models i.e. $Model_a$ & $Model_b$ are trained for class ‘a’ and ‘b’, respectively. For $Model_a$ all the reviews which belong to class ‘a’ are assigned binary class ‘1’. In contrast, reviews that do not belong to class ‘a’ are assigned binary class ‘0’. The same procedure is applied to $Model_b$ for class ‘b’. For label powerset approach, each unique combination of labels are mapped to some other unique labels. In the given example, there are three unique label combinations i.e. T_3 & T_5 has ‘a’, T_2 has ‘b’ and T_1 & T_4 has ‘a,b’. Each of these labels are mapped to some random unique classes say, ‘1’, ‘2’ & ‘3’, respectively. We use the following features for training the multi-label classifier: lexical features like n-grams, non-contiguous n-grams, character n-grams etc. For n-grams, we consider unigrams, bigrams and trigrams. Non-contiguous n-gram sequence is a pair of tokens that are n-tokens apart from each other. It helps to capture co-occurrences of terms that are far apart from each other.

	Review	$\langle T_1; T_2; T_3; T_4; T_5 \rangle$
	Label	$\langle \text{'a,b'}; \text{'b'}; \text{'a'}; \text{'a,b'}; \text{'a'} \rangle$
Binary Relevance Approach		
$Model_a$ for	Review	$\langle T_1; T_2; T_3; T_4; T_5 \rangle$
class ‘a’	Label	$\langle \text{'1'}; \text{'0'}; \text{'1'}; \text{'1'}; \text{'1'} \rangle^*$
$Model_b$ for	Review	$\langle T_1; T_2; T_3; T_4; T_5 \rangle$
class ‘b’	Label	$\langle \text{'1'}; \text{'1'}; \text{'0'}; \text{'1'}; \text{'0'} \rangle^*$
*Binary labels 1 or 0 (On or Off)		
Label Powerset Approach		
	Review	$\langle T_1; T_2; T_3; T_4; T_5 \rangle$
	Label	$\langle \text{'3'}; \text{'2'}; \text{'1'}; \text{'3'}; \text{'1'} \rangle^\wedge$
^Assign unique labels to each combination: ‘a’ =>‘1’; ‘b’ =>‘2’; ‘a,b’ =>‘3’		

Table 5: A hypothetical example for multi-label learning using binary relevance and label powerset techniques.

3.2 Sentiment Classification

Once the aspect categories are identified, we classify them to one of the four sentiment polarity classes, namely *positive*, *negative*, *neutral* and *conflict*. For each aspect category in a review we define a tuple, made up of review text and specific category, and feed it to the learning model to detect the sentiments. For e.g., if a review text ‘T’ has two aspect categories ‘food’ and ‘price’, then we define two tuples as $\langle T, \text{food} \rangle$ and $\langle T, \text{price} \rangle$ as an input to the system. Here, we use basic lexical features like n-grams, non-contiguous n-grams, character n-grams along with PoS tag and semantic orientation (SO) [18] score which is a measure of association of tokens towards negative and positive sentiments, and can be defined as:

$$SO(t) = PMI(t, posRev) - PMI(t, negRev) \quad (1)$$

where $PMI(t, negRev)$ stands for point-wise mutual information of a token t towards negative sentiment reviews. The SO score would be more effective had we use external data, but in this paper we restrict ourselves not to use any external resources for the sake of domain and resource independence.

4 Experimental Result and Analysis

To address the problem of multi-label classification of aspect category detection, we use MEKA³ for the experiments. MEKA is an extension to WEKA which handles multi-label scenario. As a base classifier we use naive Bayes [19], J48 [20] implementation of decision tree and SMO [21] implementation of SVM [22]. The underlying experiment is carried out by the following two approaches i.e. binary relevance method and label powerset method. For the label powerset approach we use MULAN⁴ [23] framework. For the sake of experiment we combine the reviews of all the electronics products, *except* mobile apps, travels and movies, and treat them as to belong to a single domain, namely '*electronics*'. Therefore, we build our model for the four major domains i.e. '*electronics*', '*mobile apps*', '*travels*' and '*movies*'. To evaluate the system we use the evaluation script, which was provided by the SemEval shared task organizer. We perform 3-fold cross validation on the training dataset. We obtain the average F-measures of 46.46%, 56.63%, 30.97% and 64.27% for aspect category detection task in electronics, mobile apps, travels and movies domain, respectively. Naive Bayes performs better in electronics and mobile apps domain, while decision tree reports better results for the travels and movies domain. In sentiment classification our proposed model reports the accuracy of 54.48%, 47.95%, 65.20% & 91.62% for the four domains respectively. Experimental results for the two tasks are reported in Table 6. We perform error analysis in order to understand the quality of the results that we obtain. An overview of the different kinds of errors encountered for aspect category detection is shown in the confusion matrix as shown in Table 7. Results show that the system obtains good recall for the '*hardware*' category, but precision is not so impressive. The model does not perform well for the other categories. One possible reason behind this could be the presence of a relatively fewer number of instances for all the domains *except* '*hardware*' which is a dominating category in the dataset i.e. 1,797 out of 3,614 instances belong to this particular category. Confusion matrix for sentiment classification is shown in Table 8. It shows that classifier performs better for the *positive* class, and this could be due to the higher number of instances, belonging to this particular class. It classifies 1,120 instances correctly out of total 1,635. The level of accuracy that we obtain for the '*neutral*' class requires further investigation. Lack of sufficient number of instances drives the system to predict only 2 correct instances for the '*conflict*' class.

³ <http://meka.sourceforge.net/>

⁴ <http://mulan.sourceforge.net/>

Domain	Method	Aspect Category Detection						Sentiment Classification
		Binary Rel.			MULAN			WEKA
		Pre	Rec	F	Pre	Rec	F	Accuracy
Electronics	NB	31.62	37.63	34.37	48.00	45.05	46.46	50.95
	DT	49.61	17.28	25.63	31.73	31.73	31.73	54.48
	SMO	26.70	46.93	34.03	39.36	44.90	41.94	51.07
Mobile Apps	NB	39.30	46.19	42.47	59.20	54.09	56.53	46.78
	DT	44.28	41.75	42.97	85.07	24.89	38.51	47.95
	SMO	51.73	38.47	44.12	45.77	57.14	50.82	42.10
Travels	NB	26.84	26.88	26.86	20.87	31.90	25.23	56.06
	DT	27.98	22.73	25.08	99.82	18.33	30.97	65.20
	SMO	25.51	20.67	22.83	15.61	39.55	22.38	60.63
Movies	NB	41.99	65.44	51.15	56.66	63.32	59.81	87.78
	DT	47.45	58.12	52.24	64.16	64.38	64.27	91.62
	SMO	43.78	59.81	50.55	48.60	63.26	54.97	91.62

Table 6: Results of aspect category detection and sentiment classification. Here, NB: naive Bayes classifier, DT: Decision tree classifier and SMO: Sequential minimal optimization implementation of SVM

5 Conclusion

In this paper we have proposed a benchmark setup for aspect category detection and its sentiment classification for Hindi. We have collected review sentences from the various online sources and annotated 5,417 review sentences across 12 domains. Based on these datasets we develop frameworks for aspect category detection and sentiment classification based on supervised classifiers. The problem of aspect category detection was cast as a multi-label classification problem whereas sentiment classification was modeled as a multi-class classification problem. The proposed model reports 46.46%, 56.63%, 30.97% and 64.27% F-measures for the aspect category detection in electronics, mobile apps, travels and movies domain, respectively. For sentiment classification the model we obtain the accuracies of 54.48%, 47.95%, 65.20% & 91.62% for the four domains, respectively. The key contributions of the research reported here are two-fold, i.e. creating a benchmark set up for aspect category detection and sentiment classification, and developing a benchmark setup that can be used as a reference for further research.

To the best of our knowledge, this is the very first attempt for these two specific problems involving Indian languages, especially Hindi. In future we would like to use domain-specific features for the problems and investigate deep learning methods for the tasks.

	Hardware	Software	Desing	Price	Ease	Misc	NoClass
Hardware	1651	0	0	0	0	0	146
Software	0	12	0	0	0	0	358
Design	0	0	39	0	0	0	485
Price	0	0	0	3	0	0	225
Ease	0	0	0	0	0	0	122
Misc	0	0	0	0	0	30	543
NoClass	1566	78	198	63	15	196	15402

Table 7: Confusion matrix for aspect category detection in electronics domain

	Positive	Negative	Neutral	Conflict
Positive	1120	77	434	4
Negative	290	96	138	0
Neutral	642	64	628	1
Conflict	73	11	34	2

Table 8: Confusion matrix for aspect category sentiment in electronics domain

References

1. Pang, B., Lee, L.: Opinion Mining and Sentiment Analysis. *Foundations and Trends in Information Retrieval* **2**(1-2) (2008) 1–135
2. Pontiki, M., Galanis, D., Pavlopoulos, J., Papageorgiou, H., Androutsopoulos, I., Manandhar, S.: Semeval-2014 task 4: Aspect based sentiment analysis. In: *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, Dublin, Ireland (August 2014)
3. Toh, Z., Wang, W.: Dlirec: Aspect term extraction and term polarity classification system. In: *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*. (2014) 235–240
4. Chernyshevich, M.: IHS R&D Belarus: Cross-domain Extraction of Product Features using Conditional Random Fields. (2014) 309–313
5. Wagner, J., Arora, P., Cortes, S., Barman, U., Bogdanova, D., Foster, J., Tounsi, L.: DCU: Aspect-based Polarity Classification for Semeval Task 4. In: *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*. (2014) 223–229
6. Castellucci, G., Filice, S., Croce, D., Basili, R.: Unitor: Aspect based sentiment analysis with structured learning. In: *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, Dublin, Ireland, Association for Computational Linguistics and Dublin City University (August 2014) 761–767
7. Gupta, D.K., Reddy, K.S., Ekbal, A.: PSO-ASent: Feature Selection using Particle Swarm Optimization for Aspect based Sentiment Analysis. In: *Natural Language Processing and Information Systems*. Springer (2015) 220–233

8. Joshi, A., Balamurali, A., Bhattacharyya, P.: A Fall-back Strategy for Sentiment Analysis in Hindi: A Case Study. *Proceedings of the 8th ICON* (2010)
9. Balamurali, A.R., Joshi, A., Bhattacharyya, P.: Cross-lingual Sentiment Analysis for Indian Languages using Linked Wordnets. In: *COLING 2012, 24th International Conference on Computational Linguistics*. (2012) 73–82
10. Balamurali, A.R., Joshi, A., Bhattacharyya, P.: Harnessing Wordnet Senses for Supervised Sentiment Classification. In: *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. (2011) 1081–1091
11. Bakliwal, A., Arora, P., Varma, V.: Hindi subjective lexicon: A lexical resource for hindi polarity classification. *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC)* (2012)
12. Mittal, N., Agarwal, B., Chouhan, G., Bania, N., Pareek, P.: Sentiment analysis of hindi review based on negation and discourse relation. In: *proceedings of International Joint Conference on Natural Language Processing*. (2013) 45–50
13. Sharma, R., Nigam, S., Jain, R.: Polarity detection movie reviews in hindi language. *CoRR abs/1409.3942* (2014)
14. Das, D., Bandyopadhyay, S.: Labeling emotion in bengali blog corpus—a fine grained tagging at sentence level. In: *Proceedings of the 8th Workshop on Asian Language Resources*. (2010) 47
15. Das, A., Bandyopadhyay, S.: Phrase-level polarity identification for bangla. *Int. J. Comput. Linguist. Appl.(IJCLA)* **1**(1-2) (2010) 169–182
16. Das, A., Bandyopadhyay, S., Gambäck, B.: Sentiment analysis: what is the end user’s requirement? In: *Proceedings of the 2nd International Conference on Web Intelligence, Mining and Semantics, ACM* (2012) 35
17. Akhtar, M.S., Ekbal, A., Bhattacharyya, P.: Aspect based sentiment analysis in hindi: Resource creation and evaluation. In: *Proceedings of the 10th edition of the Language Resources and Evaluation Conference (LREC)*. (2016)
18. Hatzivassiloglou, V., McKeown, K.R.: Predicting the Semantic Orientation of Adjectives. In: *Proceedings of the ACL/EACL*. (1997) 174–181
19. John, G.H., Langley, P.: Estimating continuous distributions in bayesian classifiers. In: *Eleventh Conference on Uncertainty in Artificial Intelligence, San Mateo, Morgan Kaufmann* (1995) 338–345
20. Quinlan, R.: *C4.5: Programs for Machine Learning*. Morgan Kaufmann Publishers, San Mateo, CA (1993)
21. Platt, J.C.: Sequential minimal optimization: A fast algorithm for training support vector machines. Technical report, *ADVANCES IN KERNEL METHODS - SUPPORT VECTOR LEARNING* (1998)
22. Cortes, C., Vapnik, V.: Support-vector networks. *Mach. Learn.* **20**(3) (September 1995) 273–297
23. Tsoumakas, G., Spyromitros-Xioufis, E., Vilcek, J., Vlahavas, I.: Mulan: A java library for multi-label learning. *Journal of Machine Learning Research* **12** (2011) 2411–2414