

Identifying argument-based relation properties in opinions

Pavithra Rajendran¹, Danushka Bollegala¹, and Simon Parsons²

¹ Department of Computer Science
University of Liverpool
United Kingdom

² Department of Informatics
Kings College London
United Kingdom

Abstract. We consider the problem of identifying properties which can relate propositions in opinionated texts that are part of an argument structure. Given a statement or a set of statements, the term *stance* refers to the overall viewpoint present, whether it is in favour or against the topic in discussion — a form of persuasion as in computational argumentation. Stance classification has become a prominent topic among NLP researchers and while stance alone cannot help in identifying the argumentative context behind a text, we show that opinions containing a stance can be classified in terms of the way the opinions are expressed: either as explicit or implicit opinions. We use three supervised methods — surface-based, embeddings-based and hybrid — and analyse the performance of these methods. We also propose three types of domain-based ontology relation, based on the implicit/explicit opinion classification, which can be used to relate appropriate propositions within an argument structure.

1 Introduction

Online e-commerce websites such as Amazon, TripAdvisor etc. encourage users to post reviews, and such activity has been growing as means of communication [10]. In fact, research [25] tells us that customer reviews have become an important factor influencing customer trust in online markets. At the same time, in the field of computational argumentation, persuasion has become a hot topic in recent work [11, 19, 22]. Persuasion in user reviews can be related to persuading the reader in favour or against the product or service. The monological structure of reviews is different from the dialogical context of most work in persuasion in the sense that reviewers independently persuade readers with viewpoints on a variety of aspects related to the product/service. As a result, opinionated texts or evaluative expressions are commonly found in reviews.

Argument mining refers to “the task of extracting arguments or argument components present in natural language texts” [15] and has also been extended to analysing these arguments using computational argumentation techniques.

There are two main ways of interpreting these arguments — (a) on an abstract level where an argument need not have any specific internal structure [7] and (b) on a structural level where the argument can have different propositions representing major premise and minor premise leading to a conclusion.

We are interested in the second of these, where we want to extract properties that relate propositions and hence help in forming argument structures. Our motivation is to classify opinions and extract relation-based properties that can help to fit these opinions within existing argument structures. Our approach fits with the idea of an argument as:

- A triple of major premise, minor premise, and conclusion
- Two or more propositions connected by an OVA+ direct inference relation [13] meaning that one proposition provides the reasoning behind the other.

though it is not limited to these representations only.

In the field of argument mining, very few works have targeted opinionated texts present in reviews, with lack of proper annotated corpora being a problem. Wyner et al. [28] proposed argumentation schemes that can identify certain properties of a product and how these properties can promote the value of the product and thus can identify arguments. Villalba et al. [26] identified evaluative expressions based on the properties of statements and they also show how RST relations can help in identifying arguments. Other work [6, 18] focused on using abstract argumentation methods for various other tasks.

Unlike previous works, we do not consider argumentation schemes or RST relations to identify arguments. Unlike the components identified using argumentation schemes, we do not focus on identifying certain properties satisfying the schemes. Instead, our methodology considers sentence level statements as opinions (propositions) and we identify relation-based properties of appropriate propositions that can lead to argument structures.

Our contributions in this paper are as follows:

- We propose a novel approach to classify sentence-level statements as explicit or implicit based on whether the stance (whether the reviewer is in favour or against the product/service) in these statements is expressed directly or not. This can help us in identifying relation-based properties of these propositions/opinions that can become a part of an argument structure.
- Using the results of our explicit/implicit opinion classification, we propose three types of domain-based ontology relation that can help in combining appropriate propositions together as part of an argument structure.
- We describe experiments on manually annotated sentence-level statements from hotel reviews using three different supervised approaches, *surface-based*, *embedding-based* and *hybrid*, with linguistic features and embeddings-based sentence features for classifying the statements as explicit or implicit opinions.
- An error analysis of the three different methods is also presented.

2 Related work

Lippi and Torroni [14] survey work on argument mining, pointing out that arguments have been extracted from a range of texts. Other works concentrate on using computational argumentation techniques [4, 9] for analysing and investigating the relations among arguments. Our contribution to this literature is to classify sentence-level statements and use this in identifying components that are part of an argument structure.

Our work is stance-based in the sense that we make use of whether a reviewer expresses their stance [21] or not. In reviews, the definition of stance refers to whether the reviewer is for or against the product/service. Such stance identification has become a prominent topic in NLP research [12, 1]. Few works in argument mining [3, 20, 23] have also expressed interest in stance classification. Our work differs from the above, since we are concerned with stance — whether the reviewer expresses it directly or not — and we use this to define three types of domain-based ontology relation which can help in identifying propositions that can be combined together within an argument structure.

Carston and Toni [5] use relation mining to identify arguments since they argue that certain objective statements can also become argumentative depending on the context. In our work, we work using opinionated texts that express stance and at present do not consider using objective statements. Instead of focussing on relations, we are interested in bridging the gap that occurs in texts that heavily depend on entities or aspects as found in reviews. Biran and Rambow [2] look for major premises as justification of claims — our work is broader in its notion of justification, including other explicit information.

3 Implicit/Explicit opinion classification

In this section, we define what we mean by *explicit* or *implicit* opinions. This is useful for relation identification tasks.

3.1 Definitions

Given a review, sentiment analysis based on *aspects* or terms specific to the particular product/service has become an important task that has also been recognised in the recent SemEval2016 conference (Task 6). We think that such aspects can also help in identifying argument components present in reviews. Furthermore, aspects can be grouped based on their common properties — for example, *location* and *service* are examples of aspects present in hotel reviews (see Table 1) — and we can exploit this structure.

For a given opinion, we define it as being explicit or implicit based on whether the reviewer expresses their view in favour or against the aspect or product/service directly or not.

We define what we mean by a stance containing opinion, an explicit opinion and an implicit opinion as follows:

Hotel	Location	Service	Room	Value	FrontDesk
hotel	location	service	bathroom	value	front desk
5/4/3/2/1 star	shop(s)	breakfast	bed	price	staff
inn	underground	restaurant	decor	cheap	receptionist
motel	transport	laundry	suite	overprice	check-in
	route	bar	internet	money	manager

Table 1. Examples of aspects (normal face) present within each category (bold face).

1. **Opinion** Any sentence-level statement that has a local sentiment that is positive or negative and talks about a certain aspect of the product/service is considered to be an opinion. The stance of the reviewer is either in favour or against, as shown by the local sentiment.
2. **Explicit opinion** Any opinion that expresses the direct stance of the reviewer, that is, whether the reviewer is in favour or against is expressed directly.
3. **Implicit opinion** Any opinion that does not express the direct stance of the reviewer, that is, the reviewer provides justification or some form of reasoning that gives indirect information about whether the reviewer is in favour or against the product/service.

3.2 Annotation

The ArguAna corpus [27] contains hotel reviews from TripAdvisor.com manually annotated through crowdsourcing. The annotation contains each sentence level statement with its local sentiment (positive or negative) and the aspects present within the statement. We collect the labelled aspects from the corpus and group them into five different aspect categories based on their common properties. These five categories are *location*, *service*, *room*, *value* and *front desk* respectively. Examples of aspects present within these categories are listed in Table 1.

A total of 1861 statements³ from ArguAna reviews were manually annotated by a single annotator. To address the reliability of the annotation, a second annotator manually annotated 48 statements and the inter-annotator agreement has a Cohen’s Kappa of 0.67. Three different sets of cues helped in the manual annotation process — *general expressive cues* where words such as *recommend*, *great* and etc. help in identifying explicitly expressed stance of the reviewer, *specific expressive cues* where expressions such as *bad room* or *fast internet* can help in identifying implicit cases since the contextual notion depends on the domain in which these expressions occur and *event-based cues* where the reviewer can describe a particular incident that indirectly expresses the stance of the reviewer. A few annotated statements are present in Table 2.

³ goo.gl/vkfNkm

Opinion	Stance	Aspect	Annotation
Great <i>hotel</i> !	direct	hotel	Explicit
don't get fool by book reviews and movies, this <i>hotel</i> is not a five star luxury experience, it dosen't even have sanitary standards!	direct and indirect	hotel	Explicit
another annoyance was the <i>internet</i> access, for which you can buy a card for 5 dollars and this is supposed to give you 25 mins of access, but if you use the card more than once, it debits an access charge and rounds minutes to the nearest five.	indirect	internet	Implicit

Table 2. Examples of opinions along with the following information: whether the stance is directly (and) or indirectly expressed, the aspect present and whether the opinion is annotated explicit or implicit.

4 Implicit/explicit opinion classification

We carried out experiments using three different feature sets. Here we describe these sets, the experiments we performed, the results we obtained and an error analysis we carried out to establish whether combining features was helpful.

4.1 Feature sets

Three different methods were investigated as features for automatically classifying opinions using a supervised approach.

Surface-based method In this method, we explore the different linguistic features present within an opinion. Basic linguistic features include:

- **Unigrams and bigrams:** Each word in the opinion and each successive pair of words present in the opinion are considered.
- **Part of Speech tags (PoS):** Each word in the opinion can be tagged with a part of speech tag and each such tag is considered. Tags of successive pairs of words are also considered as features. By using the Part of Speech tags, we can abstract the lexical features that can help in increasing the recall of the classifier. We consider three main part of speech tags namely adjective, noun and verb respectively.
- **SentiwordNet scores:** In this feature, we will encode sentiment related information using the existing SentiWordNet lexicon [8]. Here, each word in the opinion is assigned a positive, negative score and an objective score such that they sum up to 1. For each word in the opinion, the difference between the positive and negative score is obtained and the computed scores are averaged and used as a feature.

Other features include:

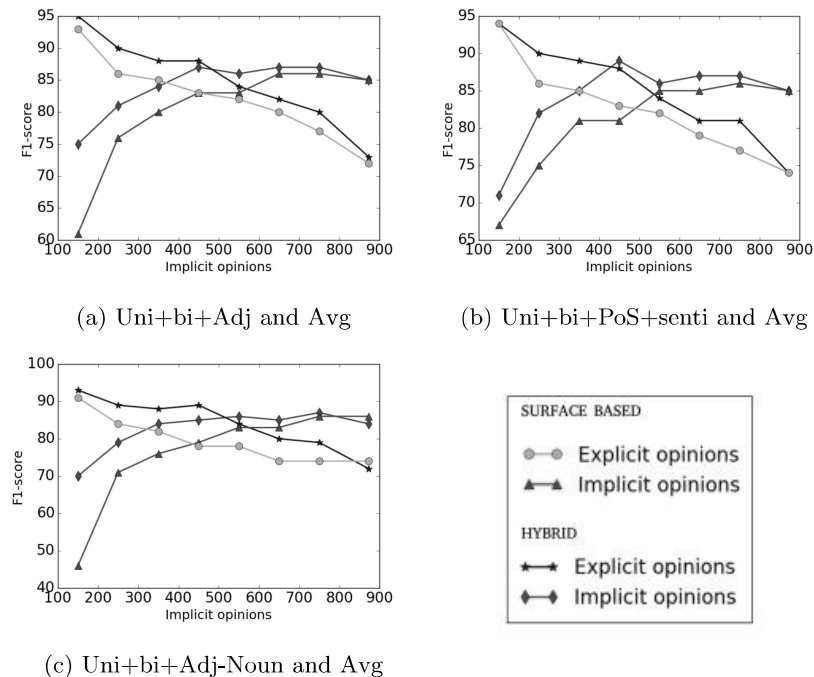


Fig. 1. Cross-validation experiments performed using surface-based features as well as hybrid-based features for different sets containing 494 explicit opinions and varying size of implicit opinions. The F1 scores are plotted against the varying implicit opinions size respectively for both the surface-based and hybrid-based methods. Three different surface-based method features using Unigrams, bigrams, PoS tags, Sentiwordnet scores and Adj-Noun pairs count are tested. In the hybrid method, we combine these three features with the average embedding-based method. Each F1-score is plotted with the corresponding marker as shown in the figure.

- **Adjective-Noun count:** For each noun present in the opinion, we combine that with the adjectives present in the same opinion and each such pair is counted as a feature.

Embedding-based method In this method, we used lower-dimensional prediction-based word embeddings as features for representing individual words, and compute a representation for an opinion as follows:

- **Average:** For each word in the opinion, the embedding vectors are collected and averaged.
- **Sum:** For each word in the opinion, the embedding vectors are collected and summation is performed.

- **Single:** For each word in the opinion, the embedding vector contains 300 different features. Each of the 300 features is considered, which means for an opinion containing n words there are $300n$ features.

Many word embeddings can be used for this purpose. In our experiments, we used 300 dimensional word embeddings⁴ pre-trained using the Global Vector Prediction (GloVe) method [16]. Each feature in the word embedding vector provides some contextual information with respect to the cooccurrences, similarity measures and etc. Here, we explore how these different features affects the classification of opinions using the above three modifications.

Hybrid method In this method, we investigate how the combination of linguistic features and embedding features and affects the classification. We combine each of the features used in the surface-based method with the *Average* embedding-based feature and use the combination as a feature. This shows how the different contextual information obtained from the embedding feature along with the linguistic feature affect the classification. The contextual knowledge captured by the embedding method can help in obtaining features that are not captured in the surface-based features.

4.2 Experiments

Data undersampling The set of 1861 annotated statements is highly imbalanced with only 494 being explicit opinions, the rest all being implicit opinions. As a result, we perform undersampling of the data using 1-Nearest Neighbour classifier [24, 17]. Initially, we keep the set of explicit opinions as training data. For every implicit opinion, if the predicted label is incorrect, then the implicit opinion along with its correct label is updated in the training data and this process continues for the rest of the implicit opinions dataset. This gave us an undersampled data containing 494 explicit opinions and 894 implicit opinions respectively.

Classifier We first carried out a comparison between different classifiers (Linear SVM, Kernel-based SVM, Logistic Regression and MultinomialNB) using unigrams and bigrams as features. This showed the linear SVM classifier outperforming the rest. Thus, we only used a linear SVM classifier in the remainder of this work. The number of explicit opinions are kept constant and the implicit opinions are varied as 150, 250, 350, 450, 550, 650, 750 and 894 respectively. A five-fold cross-validation of the different explicit-implicit sets was performed using different combination of features present in the surface-based method. These features were combined along with the average embedding-based method (this is what we call the hybrid method) and the experiment was repeated. We did not

⁴ <http://nlp.stanford.edu/data/glove.42B.300d.zip>

experiment with the other two embedding-based methods, since the average-based embedding outperformed the others. Figure 1 represents a detailed visualisation of the different F1-scores for each of the varying implicit opinions size. From the figure, it is also evident that the results improve in the case of hybrid-based method and hence features captured by the embeddings are useful in improving the overall performance.

4.3 Error analysis

We further investigated whether embedding-based features are able to capture additional contextual information that are not identified by the surface-based features. We performed an error analysis using 94 opinions from 14 different reviews. We used the 94 opinions a test set and used the rest of the opinions from the remaining reviews as the training set. The following results are reported in Table 3:

- Number of opinions correctly predicted using the surface-based method and correctly predicting using the hybrid method.
- Number of opinions correctly predicted using the surface-based method and incorrectly predicted using the hybrid method.
- Number of opinions incorrectly predicted opinions using the surface-based method and correctly predicted using the hybrid method.
- Number of opinions incorrectly predicted using the surface-based method and incorrectly predicted using the hybrid method.

and a similar analysis for the embedding-based method and the hybrid method is also reported in Table 3.

Based on the results present in the column corresponding to S_cH_c and those present in the column corresponding to E_cH_c in Table 3, we observe that the embedding method improves the performance in comparison with the surface-based method when stance is correctly predicted using the hybrid method. For instance, for a particular feature in the surface-based method, say unigrams and bigrams, we observe that there are 22 correct explicit cases in E_cH_c that outnumber the 17 correct explicit cases present in S_cH_c respectively. By comparing the results present in columns S_cH_c and $S_{ic}H_c$ with those present in columns E_cH_c and $E_{ic}H_c$ in Table 3, we can see that incorrect prediction of both the surface-based features as well as the embedding features affects the classifier performance. This shows that combination of the both these methods is better than using them separately. The rest of the results also show the embedding-based method is able to capture the features of explicit opinions better than the surface-based method.

5 Identifying relation based properties

We define three types of domain-based ontology relation which are based on the explicit/implicit classification of opinions and the aspects or aspect categories

	Type	S_cH_c	S_cH_{ic}	$S_{ic}H_c$	$S_{ic}H_{ic}$	E_cH_c	E_cH_{ic}	$E_{ic}H_c$	$E_{ic}H_{ic}$
Uni+bi	Exp	17	0	4	3	22	1	2	3
	Imp	41	13	2	4	46	0	1	8
Uni+bi+pos	Exp	21	1	2	5	21	3	2	3
	Imp	46	13	1	5	48	8	3	6
Uni+bi+senti	Exp	18	0	3	3	23	1	3	2
	Imp	43	11	2	4	46	10	1	8
Uni+bi+adj-noun	Exp	20	2	3	0	21	3	2	3
	Imp	48	9	2	5	49	7	4	5

Table 3. Error analysis of 94 opinions from 14 reviews. Opinions in each review considered as test set and the remaining as training set. Error analysis was produced based on the results for each test set or each review. S represents the surface-based method, E represents average embedding-based method and H represents the hybrid method. Subscripts c and ic indicate the number of correct and incorrect opinions. Type refers to the implicit/explicit opinion classification where exp indicates explicit and imp indicates implicit.

present in them. We define a tuple $(attribute, type, opinion)$ where $attribute$ refers to the aspect or aspect category present in an opinion denoted by $opinion$ and $type$ refers to whether it is explicit or implicit.

Definition 1. Subsumption relation

- $(attr_1, EXPLICIT, op_1)$ subsumes $(attr_2, EXPLICIT, op_2)$ if $attr_1$ is a subclass of $attr_2$, $attr_1$ is specific and $attr_2$ is generalised and must be of type ‘EXPLICIT’.

For the hotel domain, we identify aspect terms relating to the hotel as the most generalised class, followed by aspects of categories — *location, service, room, location, value* and *frontdesk*.

Definition 2. Inclusion relation Here, ‘EXPLICIT’ type is considered to be greater than ‘IMPLICIT’ since explicit opinions contain both the direct and/or the indirect stance of the reviewer, which is not the case in implicit opinions.

- $(attr_1, IMPLICIT, op_1)$ is inclusive of $(attr_2, EXPLICIT, op_2)$ if $attr_1$ and $attr_2$ belong to the same class, in this case, whether they belong to the same aspect category or not such that $attr_2$ is about the aspect category. Otherwise, $attr_1$ and $attr_2$ are about the same aspect.

Definition 3. Equivalence relation Here, if two attributes are same and contain the same type, then they are considered to be equivalent.

- $(attr_1, IMPLICIT, op_1)$ is equivalent to $(attr_2, IMPLICIT, op_2)$ if $attr_1$ and $attr_2$ are about the same aspect/aspect category.
- $(attr_3, EXPLICIT, op_3)$ is equivalent to $(attr_4, EXPLICIT, op_4)$ if $attr_3$ and $attr_4$ are about the same aspect/aspect category.

The implicit/explicit opinion classification can help in identifying propositions that can become major or minor premises leading to one of the two conclusions “I recommend the product/service” or “I do not recommend the product/service”. These relation properties can also help in identifying propositions that can be connected by a direct inference relation. However, in both these cases, the relation needs to be evaluated using NLP-based tools such as semantic similarity measures, textual entailment etc.

Let us consider the following example:

Overall rating: 1 star

1. *Another major problem was that the Hotel was simply mismanaged - although the check-in time was 3pm, the hotel refused to check us in at that time saying that the rooms were not clean* (**implicit opinion**)
2. *I would not recommend this hotel to anyone as the service is horrendous.* (**explicit opinion**)

Here, assuming that the overall star rating becomes a conclusion of the review, we find that the *implicit opinion* talks about the aspects — *hotel*, *check-in* and *room* respectively. These can be represented as (hotel, IMPLICIT, 1), (check-in, IMPLICIT, 1) and (room, IMPLICIT, 1) respectively. The *explicit opinion* talks about the following aspects — *hotel*, *service* that can be represented as (hotel, EXPLICIT, 2) and (service, EXPLICIT, 2). We extract the domain ontology relation as follows: (hotel, IMPLICIT, 1) *is inclusive of* (hotel, EXPLICIT, 2), (check-in, IMPLICIT, 1) *is inclusive of* (service, EXPLICIT, 2) and (service, EXPLICIT, 2) *subsumes* (hotel, EXPLICIT, 2). Comparing these, we say that (hotel, EXPLICIT, 2) is greater than the rest.

One way of framing the above two opinions using this domain ontology would be as follows:

Major premise I would not recommend this hotel to anyone as the service is horrendous.

Minor premise Another major problem was that the Hotel was simply mismanaged - although the check-in time was 3pm, the hotel refused to check us in at that time saying that the rooms were not clean.

Conclusion I do not recommend this hotel (1 star).

The domain ontology information can also help in identifying how these opinions are combined within an argument structure like OVA+ [13]. While this approach seems plausible, it also needs to be properly evaluated and we will investigate this as part of our future work.

6 Conclusion and future work

In this paper, we automatically classify sentence-level opinions as being implicit or explicit based on whether the opinion expresses a stance. We experiment on manually annotated opinions from a set of hotel reviews that are heavily

based upon certain aspects or entities representing the hotel domain. The explicit/implicit opinion classification uses the stance of the reviewer, that is, if the reviewer is in favour or against the product/service to identify if it is expressed directly or not. For example, *I do not recommend the hotel* expresses an opinion with direct stance about the hotel. An opinion must have a directly expressed stance to be considered explicit, otherwise it is considered to be implicit. Three different methods were investigated for automatically classifying opinions as explicit or implicit namely *surface-based*, *embedding-based* and *hybrid* methods. An error analysis of these three different methods shows that the *hybrid* method, which is a combination of linguistic features and embeddings-based sentence features help in the classification. We also propose three types of domain-based ontology relation which is based on the implicit/explicit opinion classification. We briefly describe how these can help in identifying propositions that can be combined as argument structures in reviews.

References

1. A.Walker, M., Anand, P., Abbott, R., Grant, R.: Stance classification using dialogic properties of persuasion. pp. 592–596. NAACL HLT '12 (2012)
2. Biran, O., Rambow, O.: Identifying justifications in written dialogs. In: ICSC'11. pp. 162–168 (2011)
3. Boltužić, F., Šnajder, J.: Back up your stance: Recognizing arguments in online discussions. In: ACL'14. pp. 49–58 (2014)
4. Cabrio, E., Villata, S.: Combining textual entailment and argumentation theory for supporting online debates interactions. In: ACL'12. pp. 208–212 (2012)
5. Carstens, L., Toni, F.: Towards relation based argumentation mining. In: Proceedings of the 2nd Workshop on Argumentation Mining. pp. 29–34 (2015)
6. Cocarascu, O., Toni, F.: Detecting deceptive reviews using argumentation. In: Proceedings of the 1st International Workshop on AI for Privacy and Security. pp. 1–8. PrAISE '16 (2016)
7. Dung, P.M.: On the acceptability of arguments and its fundamental role in non-monotonic reasoning, logic programming and n -person games. *Artif. Intell.* 77, 321–357 (1995)
8. Esuli, A., Sebastiani, F.: Sentiwordnet: A publicly available lexical resource for opinion mining. In: LREC'06. pp. 417–422 (2006)
9. Evripidou, V., Toni, F.: Quaestio-it.com: a social intelligent debating platform. *Journal of Decision Systems* (2014)
10. Gretzel, U., Yoo, K.H.: Use and impact of online travel reviews. In: ICTT'08. pp. 35–46 (2008)
11. Habernal, I., Gurevych, I.: Which argument is more convincing? analyzing and predicting convincingness of web arguments using bidirectional lstm. In: ACL'16. pp. 1589–1599 (2016)
12. Hasan, K.S., Ng, V.: Why are you taking this stance? identifying and classifying reasons in ideological debates. In: EMNLP'14. pp. 751–762 (2014)
13. Janier, M., Lawrence, J., Reed, C.: Ova+: an argument analysis interface. In: COMMA'14. pp. 463–464 (2014)
14. Lippi, M., Torroni, P.: Argument mining: A machine learning perspective. In: TAFAP'15. pp. 163–176 (2015)

15. Palau, R.M., Moens, M.F.: Argumentation mining: the detection, classification and structure of arguments in text. In: ICAIL'09. pp. 98–107 (2009)
16. Pennington, J., Socher, R., Manning, C.D.: Glove: Global vectors for word representation. In: EMNLP'14. pp. 1532–1543 (2014)
17. Provost, F.: Machine learning from imbalanced data sets. In: AAAI 2000 Workshop on Imbalanced Data Sets (2000)
18. Rajendran, P., Bollegala, D., Parsons, S.: Assessing weight of opinion by aggregating coalitions of arguments. In: COMMA'16. pp. 431–438 (2016)
19. Rosenfeld, A., Kraus, S.: Strategical argumentative agent for human persuasion. In: ECAI'16. pp. 320–329 (2016)
20. Sobhani, P., Inkpen, D., Matwin, S.: From argumentation mining to stance classification. In: Proceedings of the 2nd Workshop on Argumentation Mining. pp. 67–77 (2015)
21. Somasundaran, S., Wiebe, J.: Recognizing stances in ideological on-line debates. In: Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text. pp. 116–124 (2010)
22. Stab, C., Gurevych, I.: Identifying argumentative discourse structures in persuasive essays. In: EMNLP'14. pp. 46–56 (2014)
23. Toledo-Ronen, O., Bar-Haim, R., Slonim, N.: Expert stance graphs for computational argumentation. ACL'16 p. 119 (2016)
24. Tomek, I.: Two modifications of cnn. IEEE Transactions on System, Man and Cybernetics 6, 769 – 772 (1976)
25. Utz, S., Kerkhof, P., van den Bos, J.: Consumers rule: How consumer reviews influence perceived trustworthiness of online stores. Electronic Commerce Research and Applications 11, 49–58 (2012)
26. Villalba, M.P.G., Saint-Dizier, P.: A framework to extract arguments in opinion texts. IJCINI 6, 62–87 (2012)
27. Wachsmuth, H., Trenkmann, M., Stein, B., Engels, G., Palakarska, T.: A review corpus for argumentation analysis. In: CICLing'14. pp. 115–127. Springer (2014)
28. Wyner, A., Schneider, J., Atkinson, K., Bench-Capon, T.J.: Semi-automated argumentative analysis of online product reviews. In: COMMA'12. pp. 43–50 (2012)