

Argument Extraction from News, Blogs, and Social Media

Theodosios Goudas¹, Christos Louizos²,
Georgios Petasis³, and Vangelis Karkaletsis³

¹ Department of Digital Systems
University of Piraeus
Athens, Greece
`theodosios.fox@gmail.com`

² Department of Informatics & Telecommunications
University of Athens
Athens, Greece
`chr.louizos@gmail.com`

³ Software and Knowledge Engineering Laboratory,
Institute of Informatics and Telecommunications,
National Centre for Scientific Research (N.C.S.R.) “Demokritos”,
GR-153 10, P.O. BOX 60228, Aghia Paraskevi, Athens, Greece
`petasis,vangelis@iit.demokritos.gr`

Abstract. Argument extraction is the task of identifying arguments, along with their components in text. Arguments can be usually decomposed into a claim and one or more premises justifying it. Among the novel aspects of this work is the thematic domain itself which relates to Social Media, in contrast to traditional research in the area, which concentrates mainly on law documents and scientific publications. The huge increase of social media communities, along with their user tendency to debate, makes the identification of arguments in these texts a necessity. Argument extraction from Social Media is more challenging because texts may not always contain arguments, as is the case of legal documents or scientific publications usually studied. In addition, being less formal in nature, texts in Social Media may not even have proper syntax or spelling. This paper presents a two-step approach for argument extraction from social media texts. During the first step, the proposed approach tries to classify the sentences into “sentences that contain arguments” and “sentences that don’t contain arguments”. In the second step, it tries to identify the exact fragments that contain the premises from the sentences that contain arguments, by utilizing conditional random fields. The results exceed significantly the base line approach, and according to literature, are quite promising.

1 Introduction

Argumentation is a branch of philosophy that studies the act or process of forming reasons and of drawing conclusions in the context of a discussion, dialogue,

or conversation. Being an important element of human communication, its use is very frequent in texts, as a means to convey meaning to the reader. As a result, argumentation has attracted significant research focus from many disciplines, ranging from philosophy to artificial intelligence. Central to argumentation is the notion of *argument*, which according to [1] is “a set of assumptions (i.e. information from which conclusions can be drawn), together with a conclusion that can be obtained by one or more reasoning steps (i.e. steps of deduction)”. The conclusion of the argument is often called the claim, or equivalently the consequent or the conclusion of the argument, while the assumptions are called the support, or equivalently the premises of the argument, which provide the reason (or equivalently the justification) for the claim of the argument. The process of extracting conclusions/claims along with their supporting premises, both of which compose an argument, is known as argument extraction and constitutes an emerging research field.

Arguments are used in the context of a live or textual dialogue. An argument is the part of the sentence which contains one or more premises, that serve as a support to a claim, which is the conclusion [1,2,3]. According to the state of the art, there are relationships between claims and premises that existing approaches exploit in order to perform the identification of arguments in a sentence. Being an emerging research field, the existing research is rather limited and focused on specific domains such as law texts[4] and scientific publications. Social Media is a much less explored domain with only one publication related to product reviews on Amazon [5].

The difficulty of processing social media texts lies in the fact that they are expressed in an informal form, and they do not follow any formal guidelines or specific rules. Therefore, if we consider the variety of different users that publish a message and the fact that most messages are simple and informal, the probability of an argumentative sentence is rather low. Furthermore, some messages may even lack proper syntax or spelling.

Although there are a number of issues and difficulties in performing argument extraction on social media, the processing of such corpora is of great importance. Nowadays, the way that we communicate has changed. If someone wants to discuss something, or just seeks advice on a specific subject of interest, he/she just “posts” or replies to “posts” in social media, possibly providing arguments about a specific subject. It is also quite possible to post something entirely irrelevant or without any support for a possible claim. Therefore, the automated argument extraction on such corpora is extremely useful in order to acquire all the informative posts/comments (containing arguments) and discard the non-informative ones (the messages without an argument). Such a process can be extremely desirable for a wide range of applications, from supporting the decision making of a potential product buyer, who needs to decide based on product reviews from owners, to summarising discussions.

Argument extraction can also help in politics. Within the political domain it could help politicians identify the peoples’ view about their political plans, laws, etc. in order to design more efficiently their policies. Additionally, it could

help the voters in deciding which policies and political parties suit them better. Social media is a domain that contains a massive volume of information on every possible subject, from religion to health and products, and it is a prosperous place for exchanging opinions. Its nature is based on debating, so there already is plenty of useful information that waits to be identified and extracted.

However, argument extraction is not an easy task, as in many cases it is difficult even for humans to distinguish whether a part of a sentence contains an argument element or not. It may require some thought to recognize the premises and the claim, and how related they are to each other in the context of a correctly composed argument. Automatic argument extraction is a quite complex procedure, but there are a number of approaches that try to tackle this problem. Following the state of the art, our approach studies the applicability of existing approaches on the domain of social media. Following a two-step approach, we classify sentences as argumentative (containing arguments) or not, through the use of machine learning techniques, such as Logistic Regression, Random Forest, Support Vector Machines, etc. As a second step, Conditional Random Fields are employed in order to extract segments that correspond to premises in argumentative sentences.

The rest of the paper is organized as follows; Section 2 refers to the related work on argument extraction, section 3 describes the proposed methodology and the corresponding features used for our approach. Section 4 presents the experimental results and the tools we utilized and finally, section 5 concludes the paper and proposes some future directions.

2 Related Work

The area of automatic argument extraction is a relative new research field, as it has already been mentioned. One implication of this, is the absence of widely used corpora in order to comparably evaluate approaches for argument extraction. A recent and extensive survey of theories of argumentation, argumentation representations and applications targeting the social semantic web can be found in [6]. However, despite the plethora of applications targeting argumentation, almost all of them rely on manual entry of arguments by the users, and do not attempt to automatically identify and extract them from documents. Since our work is focused on automatic argument extraction, we are going to present the most influential approaches that relate to the automatic identification and extraction of argument elements from texts.

Understanding discourse relations between statements is a key factor for identifying arguments and their components in a textual document. For this reason argumentation models, as well as cue words, are employed in order to find these possible discourse relations. Most recent approaches employ machine learning and statistical approaches, usually dividing the problem as a multiple step approach. Palau et al. [4,7] methodology for extracting arguments from legal documents use this type of approach: as a first step they work at the sentence level by trying to identify possible argumentative sentences. Seeing it as a classification

task, they employ feature vectors of fixed length as a representation, containing suitable features for the selected domain. Employing different classifiers, such as maximum entropy [8], naive Bayes [9], and support vector machines [10], they comparatively evaluate their approach on the Araucaria corpus¹ and on the ECHR corpus [11], achieving an accuracy of 73% and 80% respectively. As a second step they try to identify groups of sentences that refer to the same argument, using semantic distance based on the relatedness of words contained in sentences. As a third step they detect clauses of sentences through a parsing tool, which are classified as argumentative or not with a maximum entropy classifier. Then argumentative clauses are classified into premises and claims through support vector machines. The structure of the argument is identified by employing a context-free grammar that was manually created, obtaining 60% accuracy on the ECHR corpus. Another machine learning based approach, presented in Angrosh et al.[12], employs supervised learning (conditional random fields [13]) for context identification and sentence classification of sentences in the “related work” section of research articles, based on rhetorical features extracted.

There are also approaches that employ rules in order to perform the same task. Schneider and Wyner [5,14] propose a methodology on the camera-buying domain, where the actual argument extraction is performed through the usage of a rule-based system. The system is given as input an argumentation scheme and an ontology concerning the camera and its characteristic features. These are used to define the relevant parts of the document, concerning the description of the parts of the camera. After this step is performed, the argumentation schemes are populated and along with discourse indicators and other domain specific features, the rules are constructed. An interesting aspect of this work is the fact that they applied argument extraction on product reviews in an electronic shop which is related to social media, in contrast to the majority of the work presented in the area of argument extraction which focuses on legal documents and scientific publications.

3 Proposed Approach

In order to perform argument extraction in the context of social media we followed a two-step approach. The first step includes the identification of sentences containing arguments or not. This step is necessary in order to select only the sentences that contain arguments, which constitute the input for the second step. The second step involves the usage of Conditional Random Fields (CRFs) [13] in order to identify the textual fragments that correspond to claims and premises.

3.1 Step A: Identification of Argumentative Sentences

Seeing the identification of argumentative sentences as a supervised classification task, we explored a small set of machine learning classifiers, such as Logistic

¹ http://araucaria.computing.dundee.ac.uk/doku.php#araucaria_argumentation_corpus

Regression [15], Random Forest [16], Support Vector Machines [10], Naive Bayes [9] etc. Our main research axis is not to identify the best performing machine learning algorithm for the task, but rather to study the applicability of features from the state of the art to the domain of social media. The suitability of the existing features in this domain will be evaluated and existing features will be complemented with new features that are more suitable for our domain.

The features that we have examined can be classified in two categories: features selected from the state of the art approaches, and new features that look promising for the domain of our application, which involves texts from social media. The features taken from the state of the art approaches are:

1. *Position*: this feature indicates the position of the sentence inside the text. The possible values are nominals from this set {top, top-mid, middle, middle-bot, bottom} which indicate one of the five possible positions of the sentence in the document. The motivation for this feature is to check whether the position of the sentence in the document is decisive for argument existence.
2. *Comma token number*, is the number of commas inside a sentence. This feature represents the number subordinate clauses inside a sentence, based on the idea that sentences containing argument elements may have a large number of clauses.
3. *Connective number*: is the number of connectives in the sentence, as connectives usually connect subordinate clauses. This feature is also selected based on the hypothesis that sentences containing argument elements may have a large number of clauses.
4. *Verb number*: is the number of the verbs inside a sentence, which indicates the number of periods inside a sentence.
5. *Number of verbs in passive voice*: this feature is a different version of the previous feature which takes into account the voice of the verbs, counting only the one found in passive voice.
6. *Cue words*: this feature indicates the existence and the number of cue words (also known as discourse indicators). Cue words are identified through a predefined, manually constructed, lexicon. The cue words in the lexicon are structural words which indicate the connection between periods or subordinate clauses.
7. *Domain entities number*: this feature indicates the existence and the number of entity mentions of named-entities relevant to our domain, in the context of a sentence.
8. *Adverb number*: this feature indicates the number of adverbs in the context of a sentence.
9. *Word number*: the number of words in the context of a sentence. This feature is based in the hypothesis that when we have an argument, usually, we deal with a larger sentence.
10. *Word mean length*: this is a metric of the average length (in characters) of the words in the context of a sentence.

In addition to the features found in the literature, we have examined the following set of additional/complementary features:

1. *Adjective number*: the number of adjectives in a sentence may characterize a sentence as argumentative or not. We considered the fact that usually in argumentation opinions are expressed towards an entity/claim, which are usually expressed through adjectives.
2. *Entities in previous sentences*: this feature represents the number of entities in the n^{th} previous sentence. Considering a history of $n = 5$ sentences, we obtain five features, with each one containing the number of entities in the respective sentence. These features correlate to the probability that the current sentence contains an argument element.
3. *Cumulative number of entities in previous sentences*: This feature contains the total number of entities from the previous n sentences. Considering a history of $n = 5$ we obtain four features, with each one containing the cumulative number of entities from all the previous sentences.
4. *Ratio of distributions*: we created a language model from sentences that contain argument elements and one from sentences that do not contain an argument element. The ratio between these two distributions was used as a feature. We have created three ratios of language models based on unigrams, bigrams and trigrams of words. The ratio can be described as $\frac{P(X|\text{sentence contains an argument element})}{P(X|\text{sentence does not contain an argument element})}$, where $X \in \{\text{unigrams, bigrams, trigrams}\}$.
5. *Distributions over unigrams, bigrams, trigrams of part of speech tags (POS tags)*: this feature is identical to the previous one with the exception that unigrams, bigrams and trigrams are extracted from the part of speech tags instead of words.

3.2 Step B: Extraction of Claims and Premises

Once we have identified the argumentative sentences, our approach proceeds with the extraction of the segments that represent the premises and the claims. In order to perform this task Conditional Random Fields (CRFs) [13] were employed, because it is a structured prediction algorithm, required for the task of the identification of claims and premises segments. In addition, CRFs can also take local context into consideration, which is important for the nature of this problem, as it can help maintain linguistic aspects such as the word ordering in the sentence. The features utilized in this step are: *a)* the words in these sentences, *b)* gazetteer lists of known entities for the thematic domain related to the arguments we want to extract, *c)* gazetteer lists of cue words and indicator phrases, *d)* lexica of verbs and adjectives automatically acquired using Term Frequency - Inverse Document Frequency (TF-IDF) [17] between two “documents”: The first document contained all the verbs/adjectives in an argumentative sentence whereas the second one contained the verbs/adjectives from the non-argumentative ones. The reason for restricting lexica to verbs and adjectives was the fact that premises usually contain a lot of adjectives and attribute claims through verbs.

4 Empirical Evaluation

In this section the performance of the proposed approach will be examined. The performance metrics that will be used in order to evaluate our approach is accuracy, precision, recall and F1-measure. The accuracy denotes the correctness of the prediction for the instances of both classes that are to be classified. In our case where arguments are sparse compared to the sentences that do not contain arguments, accuracy is not enough as we are mainly interested in the detection of sentences that contain arguments. Precision, recall and F1-measure can complement this task. Precision denotes how well the classifier can classify instances correctly within the performed classifications, whereas recall measures the fraction of relevant instances that are correctly retrieved from all the possible instances. F1-measure combines precision and recall as the harmonic mean.

4.1 Corpus and Preparation

All the experiments were conducted on a corpus of 204 documents collected from the social media, concerning the thematic domain of renewable energy sources. All documents are written in Greek, and originate from various sources, such as news, blogs, sites, etc. The corpus was constructed by manually filtering a larger corpus, automatically collected by performing queries on popular search engines (such as Bing²), Google Plus³, Twitter⁴, and by crawling sites from a list of sources relevant to the domain of renewable energy. The selected documents were manually annotated with domain entities and text segments that correspond to argument premises. It must be noted that claims are not expressed literally in this thematic domain, but instead they are *implied*: in this specific domain claims are not represented into documents as segments, but they are implied by the author as positive or negative views of a specific renewable energy entity or technology. Thus, in our evaluation corpus, domain entities play the role of claims, as authors argument in favor or again technologies by presenting and commenting on their various advantages or disadvantages.

The corpus has a total of 16000 sentences, where only 760 of them were annotated as containing argument elements. Related corpora that were used in the evaluation of similar approaches are the Araucaria corpus [18], which is a general corpus that has a structured set of documents in English, and the ECHR corpus [19] which is a corpus that contains annotated documents from the domain of law and legal texts, which is also in English. Unfortunately, we weren't able to gain access to any of them, limiting our ability to compare the proposed approach to the current state of the art for the English language. To our knowledge, no corpus annotated with arguments exists for the Greek language.

Our approach has been implemented within the Ellogon language engineering platform [20], as well as the Weka [21] framework. Ellogon was utilized for

² <http://www.bing.com/>

³ <https://plus.google.com/>

⁴ <https://twitter.com/>

the linguistic processing of the Greek language (tokenisation, sentence splitting, part-of-speech tagging, cue word lookup, etc.) and the creation of the feature vectors. The first step of our approach, concerning the classification of sentences as argumentative or not, was performed with the help of Weka. The second step of our approach, which is the identification of the segments of premises, was performed with the help of the CRF implementation contained in Ellogon.

4.2 Base Case

Since this specific corpus is used for the first time for argument extraction it is useful to calculate a base case that can be used to measure the performance of our approach. For this reason we have constructed a simple base case classifier: All manually annotated segments (argument components) are used in order to form a gazetteer, which is then applied on the corpus in order to detect all exact matches of all these segments. All segments identified by this gazetteer are marked as argumentative segments, while all sentences that contain at least one argumentative segment identified by the gazetteer, are characterised as an argumentative sentence. Then argumentative segments/sentences are compared to their “gold” counterparts, manually annotated by humans. Sentences that contain these recognized fragments are marked as argumentative for the first step base case, while segments marked as argumentative are evaluated for the second step base case. The results are taken through 10-fold cross validation on the whole corpus (all 16.000 sentences) and are shown in Table 1.

Table 1. Evaluation results of the base-case classifiers

	Precision	Recall	F1-Measure
Step A	14.84%	35.52%	20.50%
Step B	23.10%	21.15%	21.24%

4.3 Evaluation of the Argumentative Sentences Identification

In order to characterize and classify a sentence as a sentence which contains arguments or not, we utilized a number of well-known classifiers. Each sentence is represented by a fixed-size feature vector, using the features described in section 3, including a class representing whether it is argumentative or not. The labelled instances were used as input in order to test a variety of classifiers including Support Vector Machines, Naive Bayes, Random Forest and Logistic Regression.

The training and the evaluation of the classifiers was achieved by using the corpus already described in subsection 4.1. As already mentioned there are too many instances that correspond to sentences without arguments. So in order to create a more balanced dataset we applied a sampling which randomly ignores negative examples so as the resulting set contains an equal number of instances

from both classes. We performed two evaluations: one using 10-fold cross validation on the sampled dataset, and one splitting the initial dataset in two parts. The first part contained 70% of the instances, was sampled and used as a training set. The obtained model was evaluated on the remaining 30% of (unsampled) instances which was used as a test set. The performance of the second approach achieved 49% accuracy. The overall performance of the first approach (10-fold cross validation on sampled dataset) is shown in tables 2, 3 and 4.

Table 2. Results of various classifiers for the first step, evaluated with 10-fold cross validation (both classes)

Step A: State of the art + new features				
	Precision	Recall	F1-Measure	Accuracy
Naive Bayes	74.10%	74.00%	74.00%	73.99%
Random Forest	74.60%	74.40%	74.30%	74.38%
Logistic Regression	77.10%	77.10%	77.10%	77.12%
Support Vector Machines	76.00%	76.00%	76.00%	76.01%
Step A: State of the art features				
	Precision	Recall	F1-Measure	Accuracy
Naive Bayes	67.40%	65.40%	64.60%	65.44%
Random Forest	64.50%	64.50%	64.50%	64.47%
Logistic Regression	68.30%	68.30%	68.20%	68.25%
Support Vector Machines	68.40%	68.10%	68.00%	68.12%

Table 3. Results of various classifiers for the first step, evaluated with 10-fold cross validation (only positive class)

Step A: State of the art + new features			
	Precision	Recall	F1-Measure
Naive Bayes	72.50%	76.10%	74.30%
Random Forest	72.70%	72.50%	72.60%
Logistic Regression	76.80%	76.90%	76.80%
Support Vector Machines	74.70%	77.70%	76.20%
Step A: State of the art features			
	Precision	Recall	F1-Measure
Naive Bayes	61.30%	81.40%	69.90%
Random Forest	64.50%	62.40%	63.40%
Logistic Regression	68.60%	65.70%	67.10%
Support Vector Machines	70.30%	61.30%	65.50%

4.4 Evaluation of the Claim and Premise Segments Extraction

In order to utilize conditional random fields for the identification of premise fragments in a sentence we used the BIO representation. Each token is tagged

Table 4. Results of various classifiers for the first step, evaluated with 10-fold cross validation (only negative class)

Step A: State of the art + new features			
	Precision	Recall	F1-Measure
Naive Bayes	75.50%	71.90%	73.70%
Random Forest	73.30%	73.50%	73.40%
Logistic Regression	77.40%	77.30%	77.40%
Support Vector Machines	77.40%	74.40%	75.90%
Step A: State of the art features			
	Precision	Recall	F1-Measure
Naive Bayes	73.30%	49.90%	59.40%
Random Forest	64.50%	66.50%	65.50%
Logistic Regression	67.90%	70.80%	69.30%
Support Vector Machines	66.50%	74.80%	70.40%

Table 5. Example of the BIO representation of a sentence

BIO tag	word	prev. word	next word	...
B-premise	Wind	-	turbines	...
I-premise	turbines	Wind	generate	...
I-premise	generate	turbines	noise	...
I-premise	noise	generate	in	...
O	in	noise	the	...
O	the	in	summer	...
O	summer	the	-	...

with one of three special tags, B for starting a text segment (premise), I for a token in a premise other than the first, and O for all other tokens (outside of the premise segment). For example the BIO representation of the sentence “Wind turbines generate noise in the summer” is presented in Table 5.

The overall performance of the second step is shown in the table 6. The dataset was composed from all the sentences that contained argumentative fragments from the manual annotation.

It is clear that in both steps the results of the proposed approach are above the base case. In the first step, where we identify sentences that contain arguments, there is an increase in performance from 20% to 77%, by using the logistic regression classifier. Continuing to the second step also our results are above the base case. We have measured an increase from 22% to 43% in F1-measure, regarding the identification of the argumentative fragments, through the use of conditional random fields. Additionally, our corpus had very sparse argumentative sentences in many domains. Following the state of the art, approaches are evaluated on datasets containing an equal number of instances for argumentative and

non argumentative segments. Thus, we also performed a similar evaluation through the use of subsampling where negative examples were randomly rejected.

Table 6. Evaluation results of CRFs for the second step, evaluated with 10-fold cross validation

Step B: Identifying claim/premise segments			
	Precision	Recall	F1-Measure
CRF	62.23%	32.43%	42.37%

5 Conclusion

In this research paper we proposed a two step approach for argument extraction on a corpus obtained from social media, concerning renewable energy sources in the Greek language. In the first step we employed a statistical approach through the use of machine learning and more specifically, the logistic regression classifier. The results concerning this first phase are quite promising, since they exceeded significantly the accuracy of our base case classifier in the identification of argumentative sentences. The addition of complementary features was also justified, since they increased the performance further, thus providing a more accurate extraction of argumentative sentences. As far as the second step is concerned, CRFs are quite promising due to the fact that they are a structure prediction algorithm, required for the identification of segments, and due to their performance on the task that outperformed the base classifier.

Regarding future work, it would be interesting to explore additional new features for the first step in order to boost the accuracy even further, but considering features like verbal tense and mood, which according to [22] are good indicators of arguments. In addition, we could possibly explore more sophisticated machine learning algorithms that better suit the task of identifying sentences that contain arguments. For the argumentative segment extraction we would try other structure prediction algorithms such as Markov models and explore complementary features as well. Additionally, it would be nice to have a more comparable evaluation with the rest of the state of the art, if the Araucaria and the ECHR corpora are made publicly available again. Finally since we would prefer to work on real-world data, it would also be interesting to explore techniques that can counter the unbalanced data that are present in our dataset, without sampling.

Acknowledgments. The research leading to these results has received funding from the European Union’s Seventh Framework Programme (FP7/2007-2013) under grant agreement no 288513. For more details, please see the NOMAD project’s website, <http://www.nomad-project.eu>.

References

1. Besnard, P., Hunter, A.: *Elements of Argumentation*. MIT Press (2008)
2. Blair, J., Anthony Tindale, C.W.: *Groundwork in the Theory of Argumentation*. *Argumentation library*, vol. 21. Springer (2012)
3. Cohen, C., Copi, I.M.: *Introduction to Logic*, 11th edn. Pearson Education (2001)
4. Palau, R.M., Moens, M.F.: *Argumentation mining: the detection, classification and structure of arguments in text*. In: *ICAAIL*, pp. 98–107. ACM (2009)
5. Wyner, A., Schneider, J., Atkinson, K., Bench-Capon, T.: *Semi-automated argumentative analysis of online product reviews*. In: *Proceedings of the 4th International Conference on Computational Models of Argument, COMMA 2012* (2012)
6. Schneider, J., Groza, T., Passant, A.: *A review of argumentation for the social semantic web*. *Semantic Web* 4(2), 159–218 (2013)
7. Moens, M.F., Boiy, E., Palau, R.M., Reed, C.: *Automatic detection of arguments in legal texts*. In: *ICAAIL*, pp. 225–230. ACM (2007)
8. Berger, A.L., Pietra, V.J.D., Pietra, S.A.D.: *A maximum entropy approach to natural language processing*. *Comput. Linguist.* 22(1), 39–71 (1996)
9. Nir Friedman, D.G., Goldszmidt, M.: *Bayesian network classifiers*. *Machine Learning* 29, 131–163 (1997)
10. Cortes, C., Vapnik, V.: *Support-vector networks*. *Machine Learning* 20(3), 273–297 (1995)
11. Mochales, R., Ieven, A.: *Creating an argumentation corpus: Do theories apply to real arguments?: A case study on the legal argumentation of the echr*. In: *Proceedings of the 12th International Conference on Artificial Intelligence and Law, ICAIL 2009*, pp. 21–30. ACM, New York (2009)
12. Angrosh, M.A., Cranefield, S., Stanger, N.: *Ontology-based modelling of related work sections in research articles: Using crfs for developing semantic data based information retrieval systems*. In: *Proceedings of the 6th International Conference on Semantic Systems, I-SEMANTICS 2010*, pp. 14:1–14:10. ACM, New York (2010)
13. Lafferty, J., McCallum, A., Pereira, F.: *Conditional random fields: Probabilistic models for segmenting and labeling sequence data*. In: *Proc. 18th International Conf. on Machine Learning*, pp. 282–289. Morgan Kaufmann (2001)
14. Schneider, J., Wyner, A.: *Identifying consumers’ arguments in text*. In: Maynard, D., van Erp, M., Davis, B. (eds.) *SWAIE. CEUR Workshop Proceedings*, vol. 925, pp. 31–42. CEUR-WS.org (2012)
15. Colosimo, M.S.B.: *Logistic regression analysis for experimental determination of forming limit diagrams*. *International Journal of Machine Tools and Manufacture* 46(6), 673–682 (2006)
16. Leo, B.: *Random forests*. *Machine Learning* 45(1), 5–32 (2001)
17. Manning, C.D.: *Prabhakar Raghavan, H.S.: Introduction to Information Retrieval*. Cambridge University Press (2008)
18. Reed, C., Rowe, G.: *Araucaria: Software for argument analysis, diagramming and representation*. *International Journal of AI Tools* 14, 961–980 (2004)
19. Palau, R.M., Moens, M.F.: *Argumentation mining*. *Artif. Intell. Law* 19(1), 1–22 (2011)
20. Petasis, G., Karkaletsis, V., Paliouras, G., Androutopoulos, I., Spyropoulos, C.: *Ellogon: A new text engineering platform*. In: *Third International Conference on Language Resources and Evaluation* (2002)

21. Hall, M., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P., Witten, I.H.: The weka data mining software: An update. *SIGKDD Explorations* 11(1) (2009)
22. Florou, E., Konstantopoulos, S., Koukourikos, A., Karampiperis, P.: Argument extraction for supporting public policy formulation. In: *Proceedings of the 7th Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities*, pp. 49–54. Association for Computational Linguistics, Sofia (August 2013)