Mining Bipolar Argumentation Frameworks from natural language text

Oana Cocarascu
Imperial College London
oc511@imperial.ac.uk

Francesca Toni
Imperial College London
f.toni@imperial.ac.uk

ABSTRACT

We describe a methodology for mining topic-dependent Bipolar Argumentation Frameworks (BAFs) from natural language text. Our focus is on identifying attack and support argumentative relations between texts about the same topic, treating these texts as arguments when they are argumentatively related to other texts. We illustrate our methodology on a dataset of hotel reviews and outline some possible applications using the BAFs resulting from our methodology.

1 INTRODUCTION

Argument Mining is a relatively new research area which involves, amongst others, the automatic detection in text of arguments, argument components, and relations between arguments (see [15] for an overview). Argument Mining is a complex task because of the lack of a clear argumentative structure in free natural language text. Argument Mining can be seen as a pipeline, composed of several stages, including: identifying argumentative sentences, detecting component boundaries and argument components, and determining relations between arguments.

In this paper, we focus on identifying argumentative relations of attack and support between texts, assuming that if one text attacks/supports another, then both may be considered to be argumentative, irrespectively of their stand-alone argumentativeness. This task, referred to as *Relation-based Argument Mining* (RbAM) [6], has attracted some attention lately, seen as a stand-alone classification task to be addressed by Machine Learning (ML) techniques [2, 5, 6]. As a classification problem, RbAM can be thought of as determining a class amongst *attack*, *support*, and *neither attack nor support* for any given pair of texts, to determine which type of argumentative relation the first element of the pair is in with the second. Note that RbAM does not rely on or assume any specific argument model or internal structure of arguments.

We propose a methodology for mining Bipolar Argumentation Frameworks (BAFs) [7] from natural language text, with RbAM at its core. BAFs are a well known kind of argumentation framework in the literature on Argumentation & AI [19, 21]. RbAM is well-suited for mining BAFs since attack and support relations between (abstract) arguments are the main components of BAFs. Our methodology relies upon constructing a topic-dependent BAF from text, using topics to identify pairs of chunks of text to be classified using RbAM, along a temporal dimension whereby more recent chunks of text may either support or attack less recent ones, but not vice versa. This choice is well suited in settings such as online reviews, where information is provided incrementally, over time, and contributors have full vision of existing data.

We illustrate our methodology with natural language text drawn from a dataset of hotel reviews [17]. In this context, the topics are represented by aspects that users mention in the reviews, and BAFs represent how arguments from reviews relate to arguments from other reviews as well as to arguments about the quality of the items being reviewed. For example, consider the following two reviews about a hotel (with the second being more recent than the first):

r_1: *Exceeds any expectations - rooms, food, atmosphere were heads above anywhere. Thanks for making my trip the best of the best.*
r_2: *The room was not clean. Don't waste your time or money here.*

For the topic *room*, our methodology may give the BAF graphically shown in Figure 1 (where nodes of the graph are arguments in the BAF, edges labelled by + indicate support and edges labelled by - indicate attack). Here, the (root) G (stating that the hotel is good) is supported by G_{room} (stating that the rooms in this hotel are good), related to the topic identified. Then, α_1 and α_2, drawn from reviews r_1 and r_2 respectively, are about the same topic *room*. Here, α_1 *supports* G_{room} and α_2 *attacks* α_1. Because of the temporal dimension underlying our methodology, α_1 does not attack α_2 and α_2 does not attack G_{room}.

G: the hotel is good
$+\uparrow$
G_{room}: the rooms in this hotel are good
$+\uparrow$
α_1: rooms exceeded any expectations
$-\uparrow$
α_2: the room was not clean

Figure 1: BAF extracted from reviews r_1 and r_2.

In [9], we have shown how mining automatically BAFs from text, and reviews in particular, can be useful to support other activities (deception detection in [9]). In the current paper, we discuss some additional applications that can be supported once the BAFs have been extracted from text.

The remainder of this paper is organised as follows. We present background on BAFs and related work in Section 2. In Section 3 we outline our methodology for mining BAFs from natural language text and illustrate its application to an excerpt from a dataset of reviews about a hotel. We discuss some applications of BAFs mined from text in Section 4 and conclude in Section 5. We do not report any experimental results for our methodology in this paper; some experiments are described in [9].

2 BACKGROUND AND RELATED WORK

(Abstract) Argumentation Frameworks (AAFs) are pairs consisting of a set of arguments and a binary relation between arguments, representing attacks [11]. Formally, an AAF is any $\langle AR, attacks\rangle$ where $attacks \subseteq AR \times AR$. Bipolar Argumentation Frameworks (BAFs) extend AAFs by considering two independent binary relations between arguments: attack and support [7]. Formally, a BAF is any $\langle AR, attacks, supports\rangle$ where $\langle AR, attacks\rangle$ is an AAF and $supports \subseteq AR \times AR$. For example, consider the following three texts:

t_1: *'We should grant politicians immunity from prosecution'*
t_2: *'Giving politicians immunity from prosecution allows them to focus on performing their duties'*
t_3: *'The ability to prosecute politicians is the ultimate protection against the abuse of power'*

Here t_2 *supports* t_1 and t_3 *attacks* t_1. Thus, these three texts can be seen as the arguments in the BAF represented as the following graph (where nodes are arguments, edges labelled - indicate *attacks* and edges labelled + indicate *supports*):

$$t_2 \xrightarrow{+} t_1 \xleftarrow{-} t_3$$

Argument Mining is an emerging field whose aims include to identify argumentative sentences, argument components and argument structures (such as claims and premises), as well as to identify relations between arguments (such as support and attack) (see [15] for a recent overview). Argument Mining mostly relies on Natural Language Processing (NLP) and Machine Learning (ML) techniques. In this paper we focus on Relation-based Argument Mining (RbAM) [6], a subtask in Argument Mining which aims to automatically identify argumentative relations between texts, of the kinds (attack and support) occurring in BAFs. In RbAM, if one text attacks/supports another, then both may be considered to be argumentative, irrespectively of their stand-alone argumentativeness. For example, consider the following sentence:

Councilwoman Radcliffe voted in favour of the tax increase.

Analysed in isolation, this sentence does not seem to be argumentative but becomes an argument when read in context:

Councilwoman Radcliffe voted in favour of the tax increase. No one who voted in favour of the tax increase is a desirable candidate. Therefore, Councilwoman Radcliffe is not a desirable candidate.

RbAM can be seen as a prerequisite for constructing BAFs. It has been traditionally treated as a ML classification problem with three classes: *support*, *attack*, *neither support nor attack*.

Various approaches have been used to determine (attack/support) relations between arguments, varying from standard ML classifiers [5] to textual entailment [3, 4]. Extracting attack and support relations was also done on a corpus consisting of tweets [2]. Identifying attack and support relations between an evaluative expression and an argument was addressed in [12], on a French corpus covering domains such as hotels, restaurants, and politics.

3 A METHODOLOGY FOR MINING BAFS

RbAM is a difficult task which amounts to identifying the pieces of text between which there may be an argumentative relation as well as the type of relation between these pieces of text. When dealing with large texts, e.g. drawn from online product reviews or online debates, analysing every possible pair of pieces of texts in order to determine relations is simply not feasible.

We propose a methodology to extract BAFs from natural language text where the arguments that form the BAFs are clustered based on the topics extracted from the text being analysed. In order to construct BAFs, we use a temporal dimension to decide *which texts to compare* to determine the relation as well as to determine *the type of relation* between texts (i.e. support, attack, or neither). Thus, we construct topic-dependent BAFs by determining the relations between arguments that refer to the same topic, following the temporal order, in such a way that more recent texts can relate to less recent ones, but not vice versa. The rationale behind the topic-oriented approach is that arguments that mention different topics are highly unlikely to be related (i.e. neither argument supports nor attacks the other argument). The rationale behind the reliance on a temporal ordering is that it allows to limit the number of relations in BAFs. Our illustration demonstrates that this approach is useful and allows to generate BAFs that can be easily understood.

Our procedure for constructing a BAF from text is as follows:

(1) split the text into temporally ordered sentences; we thus assume that each argument extracted from the reviews is contained in a sentence, and that each sentence contains one or more potential argument;
(2) identify topics in texts and, for each topic, the sentences (potential arguments) related to the topic;
(3) for each topic, for each pair of sentences related to that topic, determine whether the most recent sentence *supports*, *attacks*, or *neither supports nor attacks* the less recent sentence; compare a sentence with its (temporally) closest less recent sentence first,

and compare the sentence with less recent sentences than its closest ones only if it neither supports nor attacks the closest ones;

(4) construct the BAF.

In the remainder of this section, we illustrate our methodology when applied to the reviews in Table 1, which represent an excerpt from a dataset consisting of positive and negative hotel (deceptive or truthful) reviews [17]. In the illustration, we will sometimes refer to the implementation given in [9] of our proposed methodology.

$\mathbf{R_1}$ This hotel is absolutely beautiful. Our room was gorgeous. I do have 1 complaint. The bar is very boring and the restaurant is not that great, but other than that I loved it. I would definitely stay there again.
$\mathbf{R_2}$ Very pleasant front staff, large rooms, and free Internet for those that are members of the loyalty program.
$\mathbf{R_3}$ Staff were helpful & friendly, room was huge with fantastic view of the river. Adjoining Bar/Restaurant, great food.
$\mathbf{R_4}$ I stayed here last August and I'm truly glad that I will never have to stay here again. The website does a great job of creating an illusion. The rooms are so much smaller than it seems on the website. The wireless internet is free but it is extremely slow. All and all, I would not recommend this hotel to anyone.
$\mathbf{R_5}$ It is one of the nicest hotels i have stayed at in my life, clean, comfortable and pretty. The rooms were clean and the staff is very caring.
$\mathbf{R_6}$ I have stayed in hotels all over the world, and this is probably the worst that I've ever experienced. The staff was unaccommodating, the front desk staff was condescending and not even remotely helpful. The room was not clean. Don't waste your time or your money here.
$\mathbf{R_7}$ The staff are polite and well poised. The rooms, hallways and facilities were exceptionally clean and tidy. During my stay, I stopped at their restaurant where I had one of the best American style meal in a while. Overall, this hotel is a place I would surely stay at again if given the chance to visit Chicago for a second time. It is truly exceptional.
$\mathbf{R_8}$ I loved the location and the amenities offered by this hotel. The room was charming with a window seat and a water view. Free wireless internet were a plus here. The staff was helpful and attentive. I would definitely stay here again.
$\mathbf{R_9}$ There was only one person from the staff at the front desk when we arrived, preoccupied with something on their computer so our presence was not acknowledged for several minutes. Then when we got to our room, I found it to be incredibly dusty. Overall it was a good stay but those two inconveniences made us question the amount of money we paid for it.
$\mathbf{R_{10}}$ The hotel is located in a hard to find location in Chicago, the restaurant is uncomfortably crowded, the staff is hard to reach, overall it was not a pleasant hotel stay.

Table 1: Reviews for a hotel in Chicago.

3.1 Illustration of Step 1

The first step in constructing BAFs is to split the texts analysed into sentences. This can be done for example with a pre-trained tokenizer for English. Sentences containing specific keywords such as *but*, *although*, *though*, *otherwise*, *however*, *unless*,*whereas*, are split since, in general, the phrases before and after these separators express different sentiments (e.g. *'The staff was nice but the room was messy'* results in two sentences with different sentiments).

At this step, in preparation for steps 3 and 4, we also determine the sentiment polarity of sentences. This polarity can be identified using a lexicon of frequently used adjectives in product reviews annotated with scores for sentiment polarity [10].

As an illustration of this first step, from $\mathbf{R_1}$, we identify the following pieces of texts containing potential arguments, with polarity (+ for positive and - for negative sentiment) as indicated:

room was gorgeous (+)
the bar is very boring and the restaurant is not that great (-)

whereas from $\mathbf{R_{10}}$ we identify the following potential argument, with polarity as indicated:

located in a hard to find location, the restaurant is uncomfortably
crowded, the staff is hard to reach (-)[1]

We will see next that some potential arguments can be split further.

3.2 Illustration of Step 2

In building a topic-dependent BAF from a set of reviews, we first identify 'topics' mentioned in the reviews.

Various approaches for identifying topics in text exist, ranging from associating each noun encountered in texts to a topic, to more advanced techniques related to topic modeling such as Latent Dirichlet Allocation (LDA) [1] and Non-negative Matrix Factorization (NMF) [14], able to uncover the underlying semantic structure of text by identifying topics and the words that belong to topics.

For our reviews, if we associate the nouns encountered in *at least a few reviews* to a topic, then we identify the following topics: *staff*, *room*, *internet*, *bar*, *restaurant*, *location.* The intuition here is that, in the case of online reviews, if a topic is controversial/debatable, then it will be mentioned by at least a few users (either to support the argument given initially or to attack it).

We then identify the sentences/arguments related to these topics. In the case of topics being associated to nouns, we extract the sentences that contain these specific nouns. For LDA/NMF, we extract the sentences containing any of the top words associated to the extracted topics.

For example, from $\mathbf{R_1}$, we identify the following arguments, with polarity and topics as indicated:

[1]Note that we use components of argumentative sentences to stand for the full sentences. For example, the first part of the latter potential argument stands for *"The hotel is located in a hard to find location in Chicago"*.

room: $a_{1,1}$: *room was gorgeous* (+)
bar: $a_{1,2}$: *the bar is very boring* (-)
restaurant: $a_{1,3}$: *the restaurant is not that great* (-)

whereas from $\mathbf{R_{10}}$ we identify:

location: $a_{10,1}$: *located in a hard to find location* (-)
restaurant: $a_{10,2}$: *the restaurant is uncomfortably crowded* (-)
staff: $a_{10,}$: *the staff is hard to reach* (-) [2]

3.3 Illustration of Step 3

Determining relations between sentences/arguments in any pair can be viewed as a three-class problem, with classification labels $\{attack, support, neither\}$. For this step, we can use ML classifiers to determine relations between sentences/arguments associated to topics identified in the previous step. In [9] we used Random Forests for classifying the type of relation between arguments identified at step 2, using the arguments' sentiment polarity as a feature.

3.4 Illustration of Step 4

The arguments in the BAF include a single special argument G (for 'good') as well as a special argument G_t per topic t, as already seen in Section 1. G_t stands for 'good as far as t is concerned', such that each G_t supports G.

We use a temporal approach for determining the relations between arguments related to topic t drawn from reviews and the special argument G_t. In particular, we assume that a newer argument (with respect to time) can either *support*, *attack*, or *neither support nor attack* a previous argument or G_t, but not vice versa. Note that the use of this temporal 'filter' is well suited in the context of online reviews, but it might not be applicable in other settings. If an argument a_t, related to topic t, does not support or attack another argument related to t from the same or some other review, as determined by RbAM at step 3, then a_t will either support or attack G_t, according to its polarity.

Figure 2 shows the BAF extracted from the reviews in Table 1. The extracted arguments occurring as nodes in the BAF are shown in Table 2. Note that not all text in the reviews contributes to the arguments in the BAF. For example, the first sentence from R_1: "*This hotel is absolutely beautiful*" does not represent an argument as it mentions 'hotel' and not a specific topic related to hotel.

As an illustration, consider the first three arguments about the topic *room*: $a_{1,1}$, $a_{2,2}$, $a_{3,2}$. As $a_{1,1}$ is the first argument that mentions the topic, it is connected to G_{room} according to its polarity (i.e. $a_{1,1}$ *supports* G_{room}). In this case step 3 deemed that $a_{2,2}$ neither attacks nor supports $a_{1,1}$ (we only consider relations from $a_{2,2}$ to $a_{1,1}$ as per our temporal approach). Then $a_{2,2}$ *supports* G_{room} as it has a positive polarity and there is no other less recent argument to be compared with. Here step 3 also identifies a support relation

Topic	Argument id	Argument
staff	$a_{2,1}$	very pleasant front staff
	$a_{3,1}$	staff were helpful & friendly
	$a_{5,2}$	the staff is very caring
	$a_{6,1}$	the staff was unaccommodating, the front desk staff was condescending and not even remotely helpful
	$a_{7,1}$	the staff are polite and well poised
	$a_{8,4}$	the staff was helpful and attentive
	$a_{9,1}$	there was only one person from the staff at the front desk when we arrived, preoccupied with something on their computer so our presence was not acknowledged for several minutes
	$a_{10,3}$	the staff is hard to reach
room	$a_{1,1}$	room was gorgeous
	$a_{2,2}$	large rooms
	$a_{3,2}$	room was huge with fantastic view
	$a_{4,1}$	the rooms are so much smaller than it seems on the website
	$a_{5,1}$	the rooms were clean
	$a_{6,2}$	the room was not clean
	$a_{7,2}$	the rooms were exceptionally clean and tidy
	$a_{8,2}$	the room was charming with a window seat and a water view
	$a_{9,2}$	our room, I found it to be incredibly dusty.
internet	$a_{2,3}$	free Internet for those that are members of the loyalty program
	$a_{4,2}$	the wireless internet is extremely slow
	$a_{8,3}$	free wireless internet were a plus here
bar	$a_{1,2}$	the bar is very boring
	$a_{3,3}$	adjoining Bar, great food
restaurant	$a_{1,3}$	the restaurant is not that great
	$a_{3,4}$	adjoining Restaurant, great food
	$a_{7,3}$	at their restaurant where I had one of the best American style meal in a while
	$a_{10,2}$	the restaurant is uncomfortably crowded
location	$a_{8,1}$	I loved the location
	$a_{10,1}$	located in a hard to find location

Table 2: Arguments extracted from the reviews in Table 1.

[2] In our notation, $a_{x,y}$ represents the yth argument from review x .

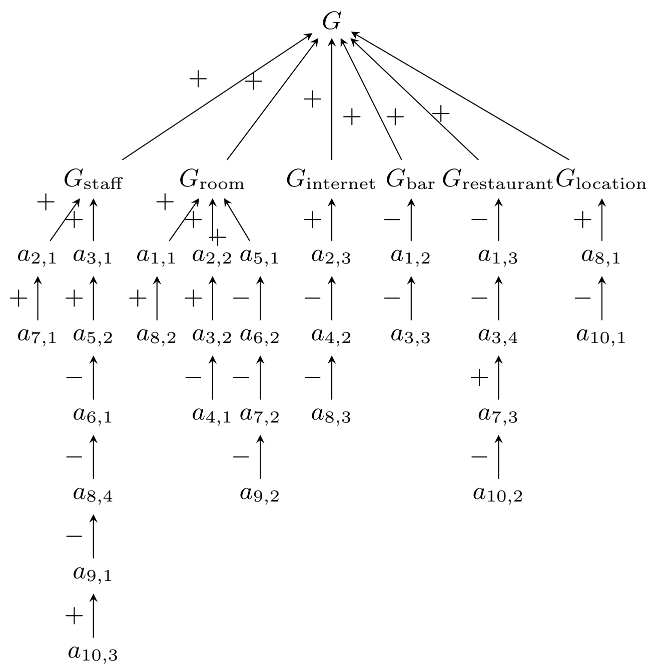

Figure 2: BAF obtained from the reviews in Table 1.

between $a_{3,2}$ and $a_{2,2}$. Thus this relation is included in the BAF.

Now consider argument $a_{3,2}$. While this can be deemed to support both $a_{2,2}$ and G_{room}, it only supports the latter in the BAF in Figure 2. Indeed, in our temporal approach we first check $a_{3,2}$ against $a_{2,2}$. If a relation is found between these arguments, then we do not check for the relation between $a_{3,2}$ and G_{room} as we want a "minimal" BAF, in terms of the number of relations it accommodates. Similarly, for $a_{4,1}$, we check for a relation between this argument and the most recent one, in this case $a_{3,2}$. If step 3 had not identified any relation between these two arguments, then $a_{4,1}$ would have been checked against $a_{2,2}$, the next "related" argument.

4 SOME APPLICATIONS

In [9] we have used the BAFs extracted from reviews to identify (argumentative) features to be fed to ML classifiers for detecting deception effectively. The BAFs provide semantic information on top of the syntactic features obtained through standard NLP techniques.

The BAFs obtained from natural language text using our methodology can be used for other purposes too. For example, if applied to online settings such as debates and reviews, various notions of dialectical strength or acceptability of arguments in AAFs and BAFs may be deployed to evaluate the outcomes of the debates or reviews, as suggested in [8]. For illustration, using the DF-QuAD method [18], that quantifies the strength of arguments by aggregating the strength of their attackers and supporters, in the case of the reviews example in Section 3, the strengths for our G_t arguments are:

$strength(G_{\text{staff}}) = 0.955078$ $strength(G_{\text{room}}) = 0.967773$
$strength(G_{\text{internet}}) = 0.6875$ $strength(G_{\text{bar}}) = 0.375$
$strength(G_{\text{restaurant}}) = 0.40625$ $strength(G_{\text{location}}) = 0.625$

These provide a measure of how good the hotel is, along the various dimensions (topics) considered, according to the available reviews, and can also be used to compare the hotel with others.

Further, BAFs could be employed in tasks such as summarisation as they provide a structured and concise view of the aspects (topics) mentioned in text. In our example hotel, from the BAF in Figure 2, we could e.g. hypothesise that *internet* used to be good, since the first review mentioning *internet* was positive ($a_{2,3}$), but has since been unstable, as the next review mentioning it is negative ($a_{4,2}$) and is followed by a positive review (see $a_{8,3}$, which attacks $a_{4,2}$).

BAFs can also help in identifying arguments that are widely accepted as well as identifying conflicting viewpoints that arise in debates. Consider the following simple example:

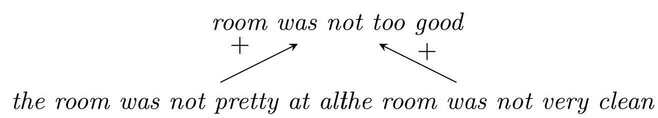

Here the reviews lead to the conclusion that the *rooms* were not good in the particular hotel under consideration, and the root argument in the graph is widely accepted.

We leave the exploration of these and additional applications of BAFs extracted by means of our methodology for future work.

5 CONCLUSION

We proposed a methodology for mining Bipolar Argumentation Frameworks (BAFs) from natural language text, relying on Relation-based Argument Mining (RbAM), a standard classification problem in NLP, to identify argumentative relations between sentences, seen as arguments by virtue of being in argumentative relations. In particular, our methodology uses RbAM to construct BAFs by determining relations between texts that refer to the same topic, along a temporal dimension whereby more recent texts may either support or attack less recent ones, but not vice versa. We have illustrated our methodology on hotel reviews and discussed the usefulness of our approach in application settings such as online user comments (reviews and debates) where arguments lack a clear structure or have incomplete/missing justifications. These applications for BAFs mined from text may help extract information and go well beyond the narrow classification task underlying standard RbAM.

This paper gives a pilot investigation, by hand, of our proposed methodology. We have referred, in our illustrations, to an implementation of our methodology [9], that also gives experimental results. Much future work is needed to explore other implementations and applicability in the settings we considered and beyond, supported by experimentation. We also plan to test whether the temporal dimension is useful in other settings, different from online reviews. We have focused on extracting BAFs from text. Other works extract different types of argument graphs (e.g. [20]), for other application areas (e.g parliamentary debates [20]). We plan to test and/or adapt our approach for this and other settings. Finally, future work also includes experimenting whether first determining arguments based on their argumentative structure, e.g. as in [13, 16], may be useful to single out chunks of text to be fed into RbAM.

REFERENCES

[1] David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent Dirichlet Allocation. *Journal of Machine Learning Research* 3 (March 2003), 993–1022.

[2] Tom Bosc, Elena Cabrio, and Serena Villata. 2016. Tweeties Squabbling: Positive and Negative Results in Applying Argument Mining on Social Media. In *Computational Models of Argument COMMA*. 21–32.

[3] Elena Cabrio and Serena Villata. 2012. Natural Language Arguments: A Combined Approach. In *ECAI*, Vol. 242. IOS Press, 205–210.

[4] Elena Cabrio and Serena Villata. 2013. Detecting Bipolar Semantic Relations among Natural Language Arguments with Textual Entailment: a Study. In *Joint Symposium on Semantic Processing (JSSP)*.

[5] Lucas Carstens. 2016. *Using Argumentation to improve classification in Natural Language problems*. Ph.D. Dissertation. Imperial College London.

[6] Lucas Carstens and Francesca Toni. 2015. Towards relation based Argumentation Mining. In *Proceedings of the 2nd Workshop on Argumentation Mining*. Association for Computational Linguistics, 29–34.

[7] Claudette Cayrol and Marie-Christine Lagasquie-Schiex. 2005. On the Acceptability of Arguments in Bipolar Argumentation Frameworks. In *Proceedings of the 8th European conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty*. Springer Berlin Heidelberg, 378–389.

[8] Federico Cerutti, Alexis Palmer, Ariel Rosenfeld, Jan Snajder, and Francesca Toni. 2016. A Pilot Study in Using Argumentation Frameworks for Online Debates. In *Proceedings of the First International Workshop on Systems and Algorithms for Formal Argumentation (SAFA)*, Vol. 1672. CEUR-WS.org, 63–74.

[9] Oana Cocarascu and Francesca Toni. 2016. Detecting deceptive reviews using Argumentation. In *Proceedings of the 1st International Workshop on AI for Privacy and Security, PrAISe@ECAI*. ACM, 9:1–9:8.

[10] Tom De Smedt and Walter Daelemans. 2012. Pattern for Python. *Journal of Machine Learning Research* 13 (2012), 2031–2035.

[11] Phan Minh Dung. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence* 77, 2 (1995), 321 – 357.

[12] Maria De La Paz Garcia Villalba and Patrick Saint-Dizier. 2013. Opinion Analysis and Argumentation: Identifying and Characterizing the Glue between Evaluative Expressions and Arguments. In *Computational Models of Natural Argumentation (CMNA)*. Proceedings of CMNA13.

[13] Ivan Habernal and Iryna Gurevych. 2017. Argumentation Mining in User-Generated Web Discourse. *Computational Linguistics* 43, 1 (2017), 125–179.

[14] Daniel D. Lee and H. Sebastian Seung. 1999. Learning the parts of objects by nonnegative matrix factorization. *Nature* 401 (1999), 788–791.

[15] Marco Lippi and Paolo Torroni. 2016. Argumentation Mining: State of the Art and Emerging Trends. *ACM Trans. Internet Techn.* 16, 2 (2016), 10.

[16] Marco Lippi and Paolo Torroni. 2016. MARGOT: A web server for argumentation mining. *Expert Systems with Applications* 65 (2016), 292 – 303.

[17] Myle Ott, Claire Cardie, and Jeffrey T. Hancock. 2013. Negative Deceptive Opinion Spam. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics.

[18] Antonio Rago, Francesca Toni, Marco Aurisicchio, and Pietro Baroni. 2016. Discontinuity-Free Decision Support with Quantitative Argumentation Debates. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Fifteenth International Conference, KR*. AAAI Press, 63–73.

[19] Iyad Rahwan and Guillermo R. Simari. 2009. *Argumentation in Artificial Intelligence* (1st ed.). Springer Publishing Company, Incorporated.

[20] Zaher Salah, Frans Coenen, and Davide Grossi. 2013. Extracting Debate Graphs from Parliamentary Transcripts: A Study Directed at UK House of Commons Debates. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Law*. 121–130.

[21] Frans H. van Eemeren, Bart Garssen, Erik C. W. Krabbe, A. Francisca Snoeck Henkemans, Bart Verheij, and Jean H. M. Wagemans. 2014. Argumentation and Artificial Intelligence. In *Handbook of Argumentation Theory*. Springer Netherlands, 615–675.