

Recognizing Stances in Ideological On-Line Debates

Swapna Somasundaran
Dept. of Computer Science
University of Pittsburgh
Pittsburgh, PA 15260
swapna@cs.pitt.edu

Janyce Wiebe
Dept. of Computer Science and
The Intelligent Systems Program
University of Pittsburgh
Pittsburgh, PA 15260
wiebe@cs.pitt.edu

Abstract

This work explores the utility of sentiment and arguing opinions for classifying stances in ideological debates. In order to capture arguing opinions in ideological stance taking, we construct an arguing lexicon automatically from a manually annotated corpus. We build supervised systems employing sentiment and arguing opinions and their targets as features. Our systems perform substantially better than a distribution-based baseline. Additionally, by employing both types of opinion features, we are able to perform better than a unigram-based system.

1 Introduction

In this work, we explore if and how ideological stances can be recognized using opinion analysis. Following (Somasundaran and Wiebe, 2009), *stance*, as used in this work, refers to an overall position held by a person toward an object, idea or proposition. For example, in a debate “Do you believe in the existence of God?,” a person may take a for-existence of God stance or an against existence of God stance. Similarly, being pro-choice, believing in creationism, and supporting universal healthcare are all examples of ideological stances.

Online web forums discussing ideological and political hot-topics are popular.¹ In this work, we are

¹<http://www.opposingviews.com>,
<http://wiki.idebate.org>, <http://www.createdebate.com> and
<http://www.forandagainst.com> are examples of such debating websites.

interested in dual-sided debates (there are two possible polarizing sides that the participants can take). For example, in a healthcare debate, participants can take a for-healthcare stance or an against-healthcare stance. Participants generally pick a side (the websites provide a way for users to tag their stance) and post an argument/justification supporting their stance.

Personal opinions are clearly important in ideological stance taking, and debate posts provide outlets for expressing them. For instance, let us consider the following snippet from a universal healthcare debate. Here the writer is expressing a negative sentiment² regarding the government (the opinion spans are highlighted in bold and their *targets*, what the opinions are about, are highlighted in italics).

- (1) *Government* is a **disease** pretending to be its own cure. [side: against healthcare]

The writer’s negative sentiment is directed toward the government, the initiator of universal healthcare. This negative opinion reveals his against-healthcare stance.

We observed that *arguing*, a less well explored type of subjectivity, is prominently manifested in ideological debates. As used in this work, *arguing* is a type of linguistic subjectivity, where a person is arguing for or against something or expressing a belief about what is true, should be true or should be done

²As used in this work, *sentiment* is a type of linguistic subjectivity, specifically positive and negative expressions of emotions, judgments, and evaluations (Wilson and Wiebe, 2005; Wilson, 2007; Somasundaran et al., 2008).

in his or her view of the world (Wilson and Wiebe, 2005; Wilson, 2007; Somasundaran et al., 2008).

For instance, let us consider the following snippet from a post supporting an against-existence of God stance.

- (2) **Obviously** that hasn't happened, and to be completely objective (as all scientists **should be**) **we must** lean on the side of **greatest evidence** which at the present time is for evolution. [side: against the existence of God]

In supporting their side, people not only express their sentiments, but they also argue about what is true (e.g., this is prominent in the existence of God debate) and about what should or should not be done (e.g., this is prominent in the healthcare debate).

In this work, we investigate whether sentiment and arguing expressions of opinion are useful for ideological stance classification. For this, we explore ways to capture relevant opinion information as machine learning features into a supervised stance classifier. While there is a large body of resources for sentiment analysis (e.g., the sentiment lexicon from (Wilson et al., 2005)), arguing analysis does not seem to have a well established lexical resource. In order to remedy this, using a simple automatic approach and a manually annotated corpus,³ we construct an arguing lexicon. We create features called *opinion-target pairs*, which encode not just the opinion information, but also what the opinion is about, its target. Systems employing sentiment-based and arguing-based features alone, or both in combination, are analyzed. We also take a qualitative look at features used by the learners to get insights about the information captured by them.

We perform experiments on four different ideological domains. Our results show that systems using both sentiment and arguing features can perform substantially better than a distribution-based baseline and marginally better than a unigram-based system. Our qualitative analysis suggests that opinion features capture more insightful information than using words alone.

The rest of this paper is organized as follows: We first describe our ideological debate data in Section 2. We explain the construction of our arguing lexicon in Section 3 and our different systems in Section

4. Experiments, results and analyses are presented in Section 5. Related work is in Section 6 and conclusions are in Section 7.

2 Ideological Debates

Political and ideological debates on hot issues are popular on the web. In this work, we analyze the following domains: Existence of God, Healthcare, Gun Rights, Gay Rights, Abortion and Creationism. Of these, we use the first two for development and the remaining four for experiments and analyses. Each domain is a political/ideological issue and has two polarizing stances: for and against.

Table 2 lists the domains, examples of debate topics within each domain, the specific sides for each debate topic, and the domain-level stances that correspond to these sides. For example, consider the Existence of God domain in Table 2. The two stances in this domain are *for*-existence of God and *against*-existence of God. “Do you believe in God”, a specific debate topic within this domain, has two sides: “Yes!!” and “No!!”. The former corresponds to the *for*-existence of God stance and the latter maps to the *against*-existence of God stance. The situation is different for the debate “God Does Not Exist”. Here, side “against” corresponds to the *for*-existence of God stance, and side “for” corresponds to the *against*-existence of God stance.

In general, we see in Table 2 that, while specific debate topics may vary, in each case the two sides for the topic correspond to the domain-level stances. We download several debates for each domain and manually map debate-level stances to the stances for the domain. Table 2 also reports the number of debates, and the total number of posts for each domain. For instance, we collect 16 different debates in the healthcare domain which gives us a total of 336 posts. All debate posts have user-reported debate-level stance tags.

2.1 Observations

Preliminary inspection of development data gave us insights which shaped our approach. We discuss some of our observations in this section.

Arguing Opinion

We found that arguing opinions are prominent when people defend their ideological stances. We

³MPQA corpus available at <http://www.cs.pitt.edu/mpqa>.

Domain/Topics	<i>stance₁</i>	<i>stance₂</i>
Healthcare (16 debates, 336 posts)	<i>for</i>	<i>against</i>
Should the US have universal health-care	Yes	No
Debate: Public insurance option in US health care	Pro	Con
Existence of God (7 debates, 486 posts)	<i>for</i>	<i>against</i>
Do you believe in God	Yes!!	No!!
God Does Not Exist	against	for
Gun Rights (18 debates, 566 posts)	<i>for</i>	<i>against</i>
Should Guns Be Illegal	against	for
Debate: Right to bear arms in the US	Yes	No
Gay Rights (15 debates, 1186 posts)	<i>for</i>	<i>against</i>
Are people born gay	Yes	No
Is homosexuality a sin	No	Yes
Abortion (13 debates, 618 posts)	<i>for</i>	<i>against</i>
Should abortion be legal	Yes	No
Should south Dakota pass the abortion ban	No	Yes
Creationism (15 debates, 729 posts)	<i>for</i>	<i>against</i>
Evolution Is A False Idea	for	against
Has evolution been scientifically proved	It has not	It has

Table 1: Examples of debate topics and their stances

saw an instance of this in Example 2, where the participant argues against the existence of God. He argues for what (he believes) is right (**should be**), and is imperative (**we must**). He employs “**Obviously**” to draw emphasis and then uses a superlative construct (**greatest**) to argue for evolution.

Example 3 below illustrates arguing in a healthcare debate. The spans **most certainly believe** and **has or must do** reveal arguing (**ESSENTIAL, IMPORTANT** are sentiments).

- (3) ... I **most certainly believe** that there are some **ESSENTIAL, IMPORTANT** things that the government **has or must do** [side: for healthcare]

Observe that the text spans revealing arguing can be a single word or multiple words. This is different from sentiment expressions that are more often single words.

Opinion Targets

As mentioned previously, a target is what an opinion is about. Targets are vital for determining

stances. Opinions by themselves may not be as informative as the *combination* of opinions and targets. For instance, in Example 1 the writer supports an against-healthcare stance using a negative sentiment. There is a negative sentiment in the example below (Example 4) too. However, in this case the writer supports a for-healthcare stance. It is by understanding *what* the opinion is about, that we can recognize the stance.

- (4) Oh, **the answer is GREEDY** insurance companies that buy your Rep & Senator. [side: for healthcare]

We also observed that targets, or in general items that participants from either side choose to speak about, by themselves may not be as informative as opinions in conjunction with the targets. For instance, Examples 1 and 3 both speak about the government but belong to opposing sides. Understanding that the former example is negative toward the government and the latter has a positive arguing about the government helps us to understand the corresponding stances.

Examples 1, 3 and 4 also illustrate that there are a variety of ways in which people support their stances. The writers express opinions about government, the initiator of healthcare and insurance companies, and the parties hurt by government run healthcare. Participants group government and healthcare as essentially the *same* concept, while they consider healthcare and insurance companies as *alternative* concepts. By expressing opinions regarding a variety of items that are same or alternative to main topic (healthcare, in these examples), they are, in effect, revealing their stance (Somasundaran et al., 2008).

3 Constructing an Arguing Lexicon

Arguing is a relatively less explored category in subjectivity. Due to this, there are no available lexicons with arguing terms (clues). However, the MPQA corpus (Version 2) is annotated with arguing subjectivity (Wilson and Wiebe, 2005; Wilson, 2007). There are two arguing categories: *positive arguing* and *negative arguing*. We use this corpus to generate a ngram (up to trigram) arguing lexicon.

The examples below illustrate MPQA arguing annotations. Examples 5 and 7 illustrate positive argu-

ing annotations and Example 6 illustrates negative arguing.

- (5) Iran **insists its nuclear program is purely for peaceful purposes.**
- (6) Officials in Panama **denied that Mr. Chavez or any of his family members had asked for asylum.**
- (7) Putin remarked that the events in Chechnia **“could be interpreted only in the context of the struggle against international terrorism.”**

Inspection of these text spans reveal that arguing annotations can be considered to be comprised of two pieces of information. The first piece of information is what we call the *arguing trigger expression*. The trigger is an indicator that an arguing is taking place, and is the primary component that anchors the arguing annotation. The second component is the expression that reveals more about the argument, and can be considered to be secondary for the purposes of detecting arguing. In Example 5, “insists”, by itself, conveys enough information to indicate that the speaker is arguing. It is quite likely that a sentence of the form “X insists Y” is going to be an arguing sentence. Thus, “insists” is an arguing trigger.

Similarly, in Example 6, we see *two* arguing triggers: “denied” and “denied that”. Each of these can independently act as arguing triggers (For example, in the constructs “X denied that Y” and “X denied Y”). Finally, in Example 7, the arguing annotation has the following independent trigger expressions “could be * only”, “could be” and “could”. The wild card in the first trigger expression indicates that there could be zero or more words in its place.

Note that MPQA annotations do not provide this primary/secondary distinction. We make this distinction to create general arguing clues such as “insist”. Table 3 lists examples of arguing annotations from the MPQA corpus and what we consider as their arguing trigger expressions.

Notice that trigger words are generally at the beginning of the annotations. Most of these are unigrams, bigrams or trigrams (though it is possible for these to be longer, as seen in Example 7). Thus, we can create a lexicon of arguing trigger expressions

Positive arguing annotations	Trigger Expr.
actually reflects Israel’s determination ...	actually
am convinced that improving ...	am convinced
bear witness that Mohamed is his ...	bear witness
can only rise to meet it by making ...	can only
has always seen usama bin ladin’s ...	has always
Negative Arguing Annotations	Trigger Expr.
certainly not a foregone conclusion	certainly not
has never been any clearer	has never
not too cool for kids	not too
rather than issuing a letter of ...	rather than
there is no explanation for	there is no

Table 2: Arguing annotations from the MPQA corpus and their corresponding trigger expressions

by extracting the starting n-grams from the MPQA annotations. The process of creating the lexicon is as follows:

1. Generate a *candidate Set* from the annotations in the corpus. Three candidates are extracted from the stemmed version of each annotation: the first word, the bigram starting at the first word, and the trigram starting at the first word. For example, if the annotation is “can only rise to meet it by making some radical changes”, the following candidates are extracted from it: “can”, “can only” and “can only rise”.
2. Remove the candidates that are present in the sentiment lexicon from (Wilson et al., 2005) (as these are already accounted for in previous research). For example, “actually”, which is a trigger word in Table 3, is a neutral subjectivity clue in the lexicon.
3. For each candidate in the *candidate Set*, find the likelihood that it is a reliable indicator of positive or negative arguing *in the MPQA corpus*. These are likelihoods of the form: $P(\text{positive arguing}|\text{candidate}) = \frac{\#\text{candidate is in a positive arguing span}}{\#\text{candidate is in the corpus}}$ and $P(\text{negative arguing}|\text{candidate}) = \frac{\#\text{candidate is in a negative arguing span}}{\#\text{candidate is in the corpus}}$
4. Make a lexicon entry for each candidate consisting of the stemmed text and the two probabilities described above.

This process results in an arguing lexicon with 3762 entries, where 3094 entries have

$P(\text{positive arguing}|\text{candidate}) > 0$; and 668 entries have $P(\text{negative arguing}|\text{candidate}) > 0$. Table 3 lists select interesting expressions from the arguing lexicon.

Entries indicative of Positive Arguing
be important to, would be better, would need to, be just the, be the true, my opinion, the contrast, show the, prove to be, only if, on the verge, ought to, be most, youve get to, render, manifestation, ironically, once and for, no surprise, overwhelming evidence, its clear, its clear that, it be evident, it be extremely, it be quite, it would therefore
Entries indicative of Negative Arguing
be not simply, simply a, but have not, can not imagine, we dont need, we can not do, threat against, ought not, nor will, never again, far from be, would never, not completely, nothing will, inaccurate and, inaccurate and, find no, no time, deny that

Table 3: Examples of positive arguing ($P(\text{positive arguing}|\text{candidate}) > P(\text{negative arguing}|\text{candidate})$) and negative arguing ($P(\text{negative arguing}|\text{candidate}) > P(\text{positive arguing}|\text{candidate})$) from the arguing lexicon

4 Features for Stance Classification

We construct opinion target pair features, which are units that capture the combined information about opinions and targets. These are encoded as binary features into a standard machine learning algorithm.

4.1 Arguing-based Features

We create arguing features primarily from our arguing lexicon. We construct additional arguing features using modal verbs and syntactic rules. The latter are motivated by the fact that modal verbs such as “must”, “should” and “ought” are clear cases of arguing, and are often involved in simple syntactic patterns with clear targets.

4.1.1 Arguing-lexicon Features

The process for creating features for a post using the arguing lexicon is simple. For each sentence in the post, we first determine if it contains a positive or negative arguing expression by looking for trigram, bigram and unigram matches (in that order) with the arguing lexicon. We prevent the same text span from matching twice – once a trigram match is found, a substring bigram (or unigram) match with the same

text span is avoided. If there are multiple arguing expression matches found within a sentence, we determine the most prominent arguing polarity by adding up the positive arguing probabilities and negative arguing probabilities (provided in the lexicon) of all the individual expressions.

Once the prominent arguing polarity is determined for a sentence, the prefix *ap* (arguing positive) or *an* (arguing negative) is attached to all the content words in that sentence to construct opinion-target features. In essence, all content words (nouns, verbs, adjectives and adverbs) in the sentence are assumed to be the target. Arguing features are denoted as *ap-target* (positive arguing toward *target*) and *an-target* (negative arguing toward *target*).

4.1.2 Modal Verb Features for Arguing

Modals words such as “must” and “should” are usually good indicators of arguing. This is a small closed set. Also, the target (what the arguing is about) is syntactically associated with the modal word, which means it can be relatively accurately extracted by using a small set of syntactic rules.

For every modal detected, three features are created by combining the modal word with its subject and object. Note that all the different modals are replaced by “should” while creating features. This helps to create more general features. For example, given a sentence “They must be available to all people”, the method creates three features “they should”, “should available” and “they should available”. These patterns are created independently of the arguing lexicon matches, and added to the feature set for the post.

4.2 Sentiment-based Features

Sentiment-based features are created independent of arguing features. In order to detect sentiment opinions, we use a sentiment lexicon (Wilson et al., 2005). In addition to positive (+) and negative (-) words, this lexicon also contains subjective words that are themselves neutral (=) with respect to polarity. Examples of neutral entries are “absolutely”, “amplify”, “believe”, and “think”.

We find the sentiment polarity of the entire sentence and assign this polarity to each content word in the sentence (denoted, for example, as *target*⁺). In order to detect the sentence polarity, we use the Vote

and Flip algorithm from Choi and Cardie (2009). This algorithm essentially counts the number of positive, negative and neutral lexicon hits in a given expression and accounts for negator words. The algorithm is used as is, except for the default polarity assignment (as we do not know the most prominent polarity in the corpus). Note that the Vote and Flip algorithm has been developed for expressions but we employ it on sentences. Once the polarity of a sentence is determined, we create sentiment features for the sentence. This is done for all sentences in the post.

5 Experiments

Experiments are carried out on debate posts from the following four domains: Gun Rights, Gay Rights, Abortion, and Creationism. For each domain, a corpus with equal class distribution is created as follows: we merge all debates and sample instances (posts) from the majority class to obtain equal numbers of instances for each stance. This gives us a total of 2232 posts in the corpus: 306 posts for the Gun Rights domain, 846 posts for the Gay Rights domain, 550 posts for the Abortion domain and 530 posts for the Creationism domain.

Our first baseline is a distribution-based baseline, which has an accuracy of 50%. We also construct *Unigram*, a system based on unigram content information, but no explicit opinion information. Unigrams are reliable for stance classification in political domains (as seen in (Lin et al., 2006; Kim and Hovy, 2007)). Intuitively, evoking a particular topic can be indicative of a stance. For example, a participant who chooses to speak about “child” and “life” in an abortion debate is more likely from an against-abortion side, while someone speaking about “woman”, “rape” and “choice” is more likely from a for-abortion stance.

We construct three systems that use opinion information: The *Sentiment* system that uses only the sentiment features described in Section 4.2, the *Arguing* system that uses only arguing features constructed in Section 4.1, and the *Arg+Sent* system that uses both sentiment and arguing features.

All systems are implemented using a standard implementation of SVM in the Weka toolkit (Hall et al., 2009). We measure performance using the accu-

racy metric.

5.1 Results

Table 4 shows the accuracy averaged over 10 fold cross-validation experiments for each domain. The first row (Overall) reports the accuracy calculated over all 2232 posts in the data.

Overall, we notice that all the supervised systems perform better than the distribution-based baseline. Observe that Unigram has a better performance than Sentiment. The good performance of Unigram indicates that what participants choose to speak about is a good indicator of ideological stance taking. This result confirms previous researchers’ intuition that, in general, political orientation is a function of “authors’ attitudes over multiple issues rather than positive or negative sentiment with respect to a single issue” (Pang and Lee, 2008). Nevertheless, the Arg+Sent system that uses both arguing and sentiment features outperforms Unigram.

We performed McNemar’s test to measure the difference in system behaviors. The test was performed on all pairs of supervised systems using all 2232 posts. The results show that there is a significant difference between the classification behavior of Unigram and Arg+Sent systems ($p < 0.05$). The difference between classifications of Unigram and Arguing approaches significance ($p < 0.1$). There is no significant difference in the behaviors of all other system pairs.

Moving on to detailed performance in each domain, we see that Unigram outperforms Sentiment for all domains. Arguing and Arg+Sent outperform Unigram for three domains (Guns, Gay Rights and Abortion), while the situation is reversed for one domain (Creationism). We carried out separate t-tests for each domain, using the results from each test fold as a data point. Our results indicate that the performance of Sentiment is significantly different from all other systems for all domains. However there is no significant difference between the performance of the remaining systems.

5.2 Analysis

On manual inspection of the top features used by the classifiers for discriminating the stances, we found that there is an overlap between the content words used by Unigram, Arg+Sent and Arguing. For

Domain (#posts)	Distribution	Unigram	Sentiment	Arguing	Arg+Sent
Overall (2232)	50	62.50	55.02	62.59	63.93
Guns Rights (306)	50	66.67	58.82	69.28	70.59
Gay Rights (846)	50	61.70	52.84	62.05	63.71
Abortion (550)	50	59.1	54.73	59.46	60.55
Creationism (530)	50	64.91	56.60	62.83	63.96

Table 4: Accuracy of the different systems

example, in the Gay Rights domain, “understand” and “equal” are amongst the top features in Unigram, while “ap-understand” (positive arguing for “understand”) and “ap-equal” are top features for Arg+Sent.

However, we believe that Arg+Sent makes finer and more insightful distinctions based on polarity of opinions toward the same set of words. Table 5 lists some interesting features in the Gay Rights domain for Unigram and Arg+Sent. Depending on whether positive or negative attribute weights were assigned by the SVM learner, the features are either indicative of for-gay rights or against-gay rights. Even though the features for Unigram are intuitive, it is not evident if a word is evoked as, for example, a pitch, concern, or denial. Also, we do not see a clear separation of the terms (for e.g., “bible” is an indicator for against-gay rights while “christianity” is an indicator for for-gay rights)

The arguing features from Arg+Sent seem to be relatively more informative – positive arguing about “christianity”, “corinthians”, “mormonism” and “bible” are all indicative of against-gay rights stance. These are indeed beliefs and concerns that shape an against-gay rights stance. On the other hand, negative arguings with these same words denote a for-gay rights stance. Presumably, these occur in refutations of the concerns influencing the opposite side. Likewise, the appeal for equal rights for gays is captured positive arguing about “liberty”, “independence”, “pursuit” and “suffrage”.

Interestingly, we found that our features also capture the ideas of opinion variety and *same* and *alternative* targets as defined in previous research (Somasingh et al., 2008) – in Table 5, items that are similar (e.g., “christianity” and “corinthians”) have similar opinions toward them for a given stance (for e.g., ap-christianity and ap-corinthians belong

to against-gay rights stance while an-christianity and an-corinthians belong to for-gay rights stance). Additionally, items that are *alternatives* (e.g. “gay” and “heterosexuality”) have opposite polarities associated with them for a given stance, that is, positive arguing for “heterosexuality” and negative arguing for “gay” reveal the the same stance.

In general, unigram features associate the choice of topics with the stances, while the arguing features can capture the concerns, defenses, appeals or denials that signify each side (though we do not explicitly encode these fine-grained distinctions in this work). Interestingly, we found that sentiment features in Arg+Sent are not as informative as the arguing features discussed above.

6 Related Work

Generally, research in identifying political viewpoints has employed information from words in the document (Malouf and Mullen, 2008; Mullen and Malouf, 2006; Grefenstette et al., 2004; Laver et al., 2003; Martin and Vanberg, 2008; Lin et al., 2006; Lin, 2006). Specifically, Lin et al. observe that people from opposing perspectives seem to use words in differing frequencies. On similar lines, Kim and Hovy (2007) use unigrams, bigrams and trigrams for election prediction from forum posts. In contrast, our work specifically employs sentiment-based and arguing-based features to perform stance classification in political debates. Our experiments are focused on determining how different opinion expressions reinforce an overall political stance. Our results indicate that while unigram information is reliable, further improvements can be achieved in certain domains using our opinion-based approach. Our work is also complementary to that by Greene and Resnik (2009), which focuses on syntactic packaging for recognizing perspectives.

For Gay Rights	Against Gay Rights
Unigram Features	
constitution, fundamental, rights, suffrage, pursuit, discrimination, government, happiness, shame, wed, gay, heterosexuality, chromosome, evolution, genetic, christianity, mormonism, corinthians, procreate, adopt	pervert, hormone, liberty, fidelity, naval, retarded, orientation, private, partner, kingdom, bible, sin, bigot
Arguing Features from Arg+Sent	
ap-constitution, ap-fundamental, ap-rights, ap-hormone, ap-liberty, ap-independence, ap-suffrage, ap-pursuit, ap-discrimination, an-government, ap-fidelity, ap-happiness, an-pervert, an-naval, an-retarded, an-orientation, an-shame, ap-private, ap-wed, ap-gay, an-heterosexuality, ap-partner, ap-chromosome, ap-evolution, ap-genetic, an-kingdom, an-christianity, an-mormonism, an-corinthians, an-bible, an-sin, an-bigot, an-procreate, ap-adopt,	an-constitution, an-fundamental, an-rights, an-hormone, an-liberty, an-independence, an-suffrage, an-pursuit, an-discrimination, ap-government, an-fidelity, an-happiness, ap-pervert, ap-naval, ap-retarded, ap-orientation, ap-shame, an-private, an-wed, an-gay, ap-heterosexuality, an-partner, an-chromosome, an-evolution, an-genetic, ap-kingdom, ap-christianity, ap-mormonism, ap-corinthians, ap-bible, ap-sin, ap-bigot, ap-procreate, an-adopt

Table 5: Examples of features associated with the stances in Gay Rights domain

Discourse-level participant relation, that is, whether participants agree/disagree has been found useful for determining political side-taking (Thomas et al., 2006; Bansal et al., 2008; Agrawal et al., 2003; Malouf and Mullen, 2008). Agreement/disagreement relations are not the main focus of our work. Other work in the area of polarizing political discourse analyze co-citations (Efron, 2004) and linking patterns (Adamic and Glance, 2005). In contrast, our focus is on document content and opinion expressions.

Somasundaran et al. (2007b) have noted the usefulness of the arguing category for opinion QA. Our tasks are different; they use arguing to retrieve relevant answers, but not distinguish stances. Our work is also different from related work in the domain of product debates (Somasundaran and Wiebe, 2009) in terms of the methodology.

Wilson (2007) manually adds positive/negative arguing information to entries in a sentiment lexicon from (Wilson et al., 2005) and uses these as arguing features. Our arguing trigger expressions are separate from the sentiment lexicon entries and are derived from a corpus. Our n-gram trigger expressions are also different from manually created regular expression-based arguing lexicon for speech data (Somasundaran et al., 2007a).

7 Conclusions

In this paper, we explore recognizing stances in ideological on-line debates. We created an arguing lex-

icon from the MPQA annotations in order to recognize arguing, a prominent type of linguistic subjectivity in ideological stance taking. We observed that opinions or targets in isolation are not as informative as their combination. Thus, we constructed opinion target pair features to capture this information.

We performed supervised learning experiments on four different domains. Our results show that both unigram-based and opinion-based systems perform better than baseline methods. We found that, even though our sentiment-based system is able to perform better than the distribution-based baseline, it does not perform at par with the unigram system. However, overall, our arguing-based system does as well as the unigram-based system, and our system that uses both arguing and sentiment features obtains further improvement. Our feature analysis suggests that arguing features are more insightful than unigram features, as they make finer distinctions that reveal the underlying ideologies.

References

- Lada A. Adamic and Natalie Glance. 2005. The political blogosphere and the 2004 u.s. election: Divided they blog. In *LinkKDD*.
- Rakesh Agrawal, Sridhar Rajagopalan, Ramakrishnan Srikant, and Yirong Xu. 2003. Mining newsgroups using networks arising from social behavior. In *WWW*.
- Mohit Bansal, Claire Cardie, and Lillian Lee. 2008. The power of negative thinking: Exploiting label disagreement in the min-cut classification framework. In

- Proceedings of the 22nd International Conference on Computational Linguistics (COLING-2008)*.
- Yejin Choi and Claire Cardie. 2009. Adapting a polarity lexicon using integer linear programming for domain-specific sentiment classification. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, pages 590–598, Singapore, August. Association for Computational Linguistics.
- Miles Efron. 2004. Cultural orientation: Classifying subjective documents by cocitation analysis. In *AAAI Fall Symposium on Style and Meaning in Language, Art, and Music*.
- Stephan Greene and Philip Resnik. 2009. More than words: Syntactic packaging and implicit sentiment. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 503–511, Boulder, Colorado, June. Association for Computational Linguistics.
- Gregory Grefenstette, Yan Qu, James G. Shanahan, and David A. Evans. 2004. Coupling niche browsers and affect analysis for an opinion mining application. In *Proceeding of RIAO-04*, Avignon, FR.
- Mark Hall, Eibe Frank, Geoffrey Holmes, Bernhard Pfahringer, Peter Reutemann, and Ian H. Witten. 2009. The weka data mining software: An update. In *SIGKDD Explorations, Volume 11, Issue 1*.
- Soo-Min Kim and Eduard Hovy. 2007. Crystal: Analyzing predictive opinions on the web. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pages 1056–1064.
- Michael Laver, Kenneth Benoit, and John Garry. 2003. Extracting policy positions from political texts using words as data. *American Political Science Review*, 97(2):311–331.
- Wei-Hao Lin, Theresa Wilson, Janyce Wiebe, and Alexander Hauptmann. 2006. Which side are you on? Identifying perspectives at the document and sentence levels. In *Proceedings of the 10th Conference on Computational Natural Language Learning (CoNLL-2006)*, pages 109–116, New York, New York.
- Wei-Hao Lin. 2006. Identifying perspectives at the document and sentence levels using statistical models. In *Proceedings of the Human Language Technology Conference of the NAACL, Companion Volume: Doctoral Consortium*, pages 227–230, New York City, USA, June. Association for Computational Linguistics.
- Robert Malouf and Tony Mullen. 2008. Taking sides: Graph-based user classification for informal online political discourse. *Internet Research*, 18(2).
- Lanny W. Martin and Georg Vanberg. 2008. A robust transformation procedure for interpreting political text. *Political Analysis*, 16(1):93–100.
- Tony Mullen and Robert Malouf. 2006. A preliminary investigation into sentiment analysis of informal political discourse. In *AAAI 2006 Spring Symposium on Computational Approaches to Analysing Weblogs (AAAI-CAAW 2006)*.
- Bo Pang and Lillian Lee. 2008. Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval*, Vol. 2(1-2):pp. 1–135.
- Swapna Somasundaran and Janyce Wiebe. 2009. Recognizing stances in online debates. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 226–234, Suntec, Singapore, August. Association for Computational Linguistics.
- Swapna Somasundaran, Josef Ruppenhofer, and Janyce Wiebe. 2007a. Detecting arguing and sentiment in meetings. In *SIGdial Workshop on Discourse and Dialogue*, Antwerp, Belgium, September.
- Swapna Somasundaran, Theresa Wilson, Janyce Wiebe, and Veselin Stoyanov. 2007b. Qa with attitude: Exploiting opinion type analysis for improving question answering in on-line discussions and the news. In *International Conference on Weblogs and Social Media*, Boulder, CO.
- Swapna Somasundaran, Janyce Wiebe, and Josef Ruppenhofer. 2008. Discourse level opinion interpretation. In *Proceedings of the 22nd International Conference on Computational Linguistics (Coling 2008)*, pages 801–808, Manchester, UK, August.
- Matt Thomas, Bo Pang, and Lillian Lee. 2006. Get out the vote: Determining support or opposition from congressional floor-debate transcripts. In *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing*, pages 327–335, Sydney, Australia, July. Association for Computational Linguistics.
- Theresa Wilson and Janyce Wiebe. 2005. Annotating attributions and private states. In *Proceedings of ACL Workshop on Frontiers in Corpus Annotation II: Pie in the Sky*.
- Theresa Wilson, Janyce Wiebe, and Paul Hoffmann. 2005. Recognizing contextual polarity in phrase-level sentiment analysis. In *hltemnlp2005*, pages 347–354, Vancouver, Canada.
- Theresa Wilson. 2007. *Fine-grained Subjectivity and Sentiment Analysis: Recognizing the Intensity, Polarity, and Attitudes of private states*. Ph.D. thesis, Intelligent Systems Program, University of Pittsburgh.