

The Politics of Comments: Predicting Political Orientation of News Stories with Commenters' Sentiment Patterns

Souneil Park¹, Minsam Ko², Jungwoo Kim¹, Ying Liu², Junehwa Song¹

¹ Department of Computer Science
KAIST, Daejeon, 305-701, Republic of Korea
{spark, jwkim, junesong}@nclab.kaist.ac.kr

² Department of Knowledge Service Engineering
KAIST, Daejeon, 305-701, Republic of Korea
{msko,yingliu}@kaist.ac.kr

ABSTRACT

Political views frequently conflict in the coverage of contentious political issues, potentially causing serious social problems. We present a novel social annotation analysis approach for identification of news articles' political orientation. The approach focuses on the behavior of individual commenters. It uncovers commenters' *sentiment patterns* towards political news articles, and predicts the political orientation from the sentiments expressed in the comments. It takes advantage of commenters' participation as well as their knowledge and intelligence condensed in the sentiment of comments, thereby greatly reduces the high complexity of political view identification. We conduct extensive study on commenters' behaviors, and discover predictive commenters showing a high degree of regularity in their sentiment patterns. We develop and evaluate sentiment pattern-based methods for political view identification.

Author Keywords

Political News, Political View Identification, Comments, Commenters, Sentiment Pattern.

ACM Classification Keywords

H5.3. Group and Organization Interfaces: Web-based interaction.

General Terms

Human Factors, Experimentation.

INTRODUCTION

The political nature of the news media often turns journalism into subjective propaganda. News producers actively frame reality in favor of their political, ideological views. They select different facets of the reality, cover these facets in different tones, and present them in different styles. Political views frequently conflict in the coverage of contentious political issues [10]. Such conflict often causes a sharp increase in political polarization [3], misunderstanding of critical issues, and a significant impact on elections [4]. However, it is difficult for ordinary readers to identify and

critically analyze the political views for a large number of news articles; they mostly read news passively. New delivery models are sought after to give readers awareness on conflicting political views.

In this paper, we present a novel social annotation analysis approach for identification of news articles' political orientation. Specifically, the approach uncovers commenters' *sentiment patterns* towards political news articles. The political orientation of articles is then predicted from the sentiments expressed in their comments. As such, it effectively deals with diverse political news articles without complicated analysis of news texts. The approach has clear advantage over the existing ones, i.e., news text analysis, and meta-data analysis. It is difficult to interpret the political orientation from computational analysis of news text. News articles cover complex political discourse; the discourse ranges over various domains such as party, government, economy, environment, etc.; different political topics involve diverse arguments. The works using this approach have been confined to several selected topics [12]. The meta-data of news articles, typically, the news company information, can be used to some extent. However, the usefulness is limited to some articles of the news companies which reveal a clear political preference [8].

The intuition behind our approach is that there exist commenters with a clear political preference and that they are highly likely to show their views consistently towards various political issues. They would actively express their preferences in a certain pattern responding to the political news articles. Based on the intuition, we discover and utilize predictive commenters, showing a high degree of regularity in their sentiment patterns; for example, a liberal predictive commenter mostly leaves a negative comment to conservative articles or a positive comment to liberal articles. The sentiment expressed by the predictive commenters strongly indicates the political orientation of the article. When the comment is negative, the article's political orientation can be predicted to be the opposite from that of the commenters; when the comment is positive, it can be predicted to be the same as that of the commenter.

The use of these commenters greatly reduces the high complexity of political view identification. Admitting the practical limitation of computational news article analysis, it takes advantage of commenters' participation as well as their

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CSCW 2011, March 19–23, 2011, Hangzhou, China.

Copyright 2011 ACM 978-1-4503-0556-3/11/03...\$10.00.

knowledge and intelligence; their interpretation of the political orientation is condensed in the sentiment of comments. Identifying sentiment from comments is relatively less complicated, specifically when interests are narrowed down to positive and negative. Comments are usually concise, often explicitly express sentiments, and frequently include explicit words such as “great” or “worse” which can be used as cues for identification.

Our approach importantly considers the behavior of individual annotators. Social annotations of news articles have been increasingly used to improve readers’ browsing experience; e.g., the number of votes or links is a commonly used measure for recommending popular news articles. Yet, most applications of social annotation are limited to simple aggregation, for instance, counting the number of annotations. They overlook important contextual information of individual annotators, e.g., topical or political preference of annotators, annotation patterns, etc. Understanding individual annotators’ behavior can potentially create new delivery models and support diverse navigational needs of readers.

To develop the proposed approach, we conduct extensive study on commenters’ behaviors. We select commenters and collect their comment history from Naver News, a popular Internet news portal in South Korea. The study results meet the three prerequisites for our method to work. First, there are active commenters who continuously comment on a large amount of articles. Second, most of them have a clear political preference either as liberal or conservative. Third, among them, there are predictive commenters.

Based on the study, we develop and evaluate sentiment pattern-based methods for political view identification. The method uses the commenters selected for the study. We develop two methods: single-commenter-based prediction, which utilizes only one commenter’s sentiment pattern; and multi commenter-based prediction, which aggregates the sentiment pattern of multiple commenters who comment on the same article. These methods are supervised techniques; however, we experimentally show that a small training set can be sufficient.

BACKGROUND AND RELATED WORK

Study on Political Views of the Internet

Recent research has increasingly investigated the political views of the Internet. The research can be viewed according to the media they deal with, i.e., social media and traditional news media. Much research has centered on social media, such as blogs [16] and social recommendation systems [15]. Example research efforts include studies on the political orientation of bloggers [9] and blog posts [16], linking activity of political blogger groups [1], and political opinions in social recommendation systems [15]. In social media, the political views are expressed freely and explicitly. The producers of social media are not restrained by journalism values such as fairness or balance, and do not go through a formal editorial process. Different from these works, our

work studies the political views of the traditional mainstream media.

It becomes more difficult to identify political views from traditional news media contents. Although they are not free from political bias [18], the news producers usually do not explicitly present their political views in the news. Many news articles mainly deliver facts than opinions except a few types of articles, e.g., editorials or columns. The political views can be also expressed in many different ways [2]: through fact selection, by omitting detailed facts or selecting information sources, e.g., stakeholders, experts, or think tanks; through writing style, by choice of words and tone; through presentation style, e.g., selection of photos. We speculate on the potential approaches for political view identification of news articles and review the existing research efforts.

We classify the potential approaches to news text analysis, meta-data analysis, and social annotation analysis. A few works have made efforts to identify political orientation through news text analysis [12]. Yet, these works are confined to several selected topics. They choose a certain issue, such as the Iraq war, and develop supervised classifiers with a sample of articles covering the issue. Applying this approach to general news articles is a very costly process. It requires continuous monitoring of new topics, and development of new classifiers for each topic. Frame analysis research of mass communication literature also analyzes news text [21]. These efforts heavily rely on manual analysis hence also cover only a few news events.

Park et al. [18][19] have attempted to classify news articles based on the covered aspects. The method is limited to specific type of articles, i.e., straight news articles, and classifies articles in an unsupervised manner, thus, it is not able to identify specific political orientations.

The meta-data analysis approach uses the data attached to news article, e.g., published date, name of the journalist, news company information. Typically, the news company information is considered to reveal the political view expressed in an article. The approach is frequently taken by media analysis research of the mass communication literature [8]. They select news companies who reveal strong political preference, and compare the difference of coverage for certain issues, e.g. health-care reform debate. Their investigation is limited to a few news companies with strong political preference.

The news company information is insufficient to solely determine the political orientation of the articles. There are many news companies whose political preference is unclear. We conducted a quick analysis, comparing the amount of articles published by the Korean news companies with strong political preferences and those with unclear political preferences. The selection of news companies with strong political preference was made by following a number of media analysis researches [8]. Other companies were classified to show unclear preference. In our analysis 70% of

the articles were published from the companies showing unclear preference.

In addition, even the news companies with strong political preference do not always publish articles supporting their views. It is frequently observed that news producers with contrasting political preference cover similar stories; for example, the coverage of the topic ‘health care reform bill’ commonly includes basic stories, such as a summary of the bill, and the senate’s reaction, regardless of news companies’ different views to the bill.

Social annotation analysis approach utilizes the data created by users such as comments, links, and votes. Social annotations are a valuable source of information for political view identification. They contain collective knowledge of readers, who have the potential to interpret the political views expressed in the articles. Although their goals are different from our work, a number of works also note the importance of the social annotations. BLEWS [6] does not approach to identify the political view expressed in news articles, but attempts to provide more contexts to them, e.g., how many liberals or conservatives are linking to the articles. It uses the links made from well-known liberal or conservative blogs. Munson et al.’s work [15] aims to increase diversity in social recommender systems. They used users’ voting records to distinguish different interests between them. They proposed an algorithm which promotes the recommendation items which win votes from users with different interests, instead of simply ranking based on the number of votes.

Study on Comments

Research has been conducted to study various aspects of comments: the volume of comments [13], relation with popularity [24], and their effect on search [25]. Many studies have discussed the value of comments as an indicator of popularity. They analyzed the relationship between the number of comments and the popularity of various blog posts [24], blogs [13], and news stories [22]. A number of works attempted to utilize comments to improve search results [25]. Gilbert et al. [7] measured that many comments of blogs tend to agree with the blogger’s opinion. These works take an aggregated approach to comments and do not focus on individual comments or commenters. In contrast, we delve into individual commenters’ behavioral characteristics and gain semantic knowledge from their comments.

Opinion Mining and Sentiment Analysis

There is a growing body of work on opinion mining and sentiment analysis [17]. It aims to automatically identify and classify opinions expressed in text documents. Compared to the political view identification problem, the works made the problem simpler by making a number of assumptions. They frequently assume that the input documents commonly express opinions towards a single item or topic, e.g., a movie. Many of them also primarily deal with documents which explicitly reveal opinions on the selected topics, e.g., movie reviews. In addition, most works apply one static classification frame, positive vs. negative, to the topic.

The political view identification problem is different from current opinion mining tasks. First, a political view identification method should deal with the articles covering various different topics. For example, some articles may cover liberals’ argument on global warming whereas others may cover conservatives’ argument on tax cuts. Even if the articles cover the same topic, they may cover different aspects of it. For the health care bill, some articles may cover the enlarged coverage whereas others may cover the expected increase of insurance premiums.

Second, the frames of argument are diverse, and more complex than the frame positive vs. negative. Political issues range over diverse areas, for example, government, foreign policies, business, labor, religion, etc. The conservative vs. liberal frame is realized in various forms depending on the issue at hand, e.g., government regulations vs. free market, employment stability vs. employment flexibility, pro-life vs. pro-choice, gun control vs. gun right, etc.

In addition, political orientation can be communicated without explicit expression of opinion or sentiment. It is often conveyed through objective sentences that include carefully selected facts. For example, a news article can cast a negative light on a government program simply by covering the increase of deficit caused by it.

OBSERVATIONS ON COMMENTERS

It is necessary to understand commenters’ behaviors in order to speculate about their potential for political view identification. Commenters could be effectively utilized if they comment on many political articles, and express different sentiments depending on the political orientation of the articles.

In this section, we report our findings on commenters’ behaviors. The commenters’ behaviors are analyzed from various perspectives. We first explore where there exists a distinctive group of active commenters who can be used for political view identification. We then identified the commenters’ political stance from their comments. For those who show a clear political orientation, we observed whether they show a sentiment pattern useful for predicting the political orientation of news articles.

Commenter Sample

The commenters were sampled from Naver News (<http://news.naver.com>). Naver News is one of the most popular news portals in South Korea, averaging 3.2 million page views per day.

We sampled commenters from two article sets with different characteristics, i.e., Popular Set and General Set. The Popular Set is composed of a collection of the ‘20 most read political news articles of the day’ for a 6 month period. As the stories are popular, they have many comments. The General Set is sampled from the Naver political issue directory. The directory records major political issues and stores about a hundred articles under each issue. We sampled all the articles of 11 issues that were updated from 2008.12 to 2009.11. The

set includes both articles with many comments and those with few.

We designated the Naver IDs as the identifier of commenters. Naver News does not allow anonymous comments; only users logged in with their own ID are permitted to leave comments. We then collected the selected ID's comments on political news articles during the publication time period. The comments responding to other comments were filtered out. We could identify these comments as Naver News supports threaded comments.

	General Set	Popular Set
# of articles	1395	3385
Publication time period	336 days (2008.12.15~2009.11.16)	184 days (2009.4.30~2009.10.31)
# of comments (avg. per article)	13340 (9.56)	140559 (41.52)
# of articles having more than 3 comments	371 (26.5%)	3019 (89.2%)

Table 1. Details of the Two Article Sets

Active Commenters

Wide Article Coverage of Active Commenters

We observed the existence of active commenters, who leave comments on a large proportion of articles. This observation shows the potential to deal with a large number of articles through only a small number of active commenters. Figure 1 shows the proportion of articles commented upon by each commenter. On the x-axis, the commenters are ranked according to the proportion. It shows that a small number of high-ranking commenters leave many more comments than ordinary commenters. In the Popular Set, the most active commenter commented on 11% of the articles while average commenters commented on less than 5% of the articles. In the General Set, the most active commenter commented on 5% of the articles while the average commenters commented on less than 1%. The numbers are lower because the General Set includes many more articles with few comments in comparison to the Popular Set.




Figure 1. Distribution of the amount of commented articles

From each set, we selected the 50 commenters who left the most comments, resulting in a total of 100 commenters. Figure 2 shows the number of articles commented upon by the top 50 commenters in each article set. The top 50 commenters of the Popular Set commented on nearly 80% of the articles. The top 50 commenters of the General Set commented on 20% of the articles. The proportion is relatively low compared to the Popular Set. However, when considering only the articles with more than 5 comments,

most of them (80.2%) were covered by the top 50 commenters. The articles with many comments are those which were more visible at the time of publication, hence, likely to be read. This indicates that these 50 commenters cover most of the articles that people are likely to read.




Figure 2. Amount of articles commented by the top 50 commenters of each set

Continuity of Participation

In both the Popular Set and General Set, the majority of the top 50 commenters continually left comments for a long period of time. They are a stable, long-term source of comments, covering diverse political articles over time.

For each commenter, we measured the duration of participation, i.e., the time gap between the first comment and the last comment (See Figure 3). In the Popular Set, 100% of commenters left comments for at least 5 months; of those commenters, 74% continued to be active even after six months. In the General Set, 90% of commenters left comments for at least 3 months, and 76% of commenters were active for six months or longer.




Figure 3. Duration of participation.




Figure 4. Frequency of comment.

The commenters also left comments frequently during their respective periods of participation. In order to assess commenting frequency, we measured the time gap between individual commenter's consecutive comments (See Figure 4). Of these time gaps, 92% were less than three days for the Popular Set. That is to say, there was a 92% chance that after making one comment, that commenter would leave another comment within three days. In the case of the General Set, the majority of the time gaps were also less than three days.

The time gaps of the General Set are relatively long due to the fewer number of comments in the set.

Commenter’s Political Orientation

We also explore whether the commenters’ political stance can clearly be identified from their comments. This would be possible if each commenter consistently shows a particular political stance throughout his or her commenting history.

Consistency of Political Orientation

We analyzed the comments of the 100 active commenters to the articles in the General and Popular sets. Both article sets covered timely, major political issues. Many issues were divisive between liberals and conservatives. For example, issues reported in the sets included the nomination of the prime minister, made by the current conservative administration. Liberals generally criticized the nominee and opposed the nomination while conservatives supported the nominee. Such issues were effective for categorizing commenters’ political orientations because they have disparate and clear liberal and conservative stances.

	General Set	Popular Set
Consistent	92% (Lib:54%, Cons: 38%)	86% (Lib:24%, Cons. 62%)
Vague	8%	14%

Table 2. The Consistency of Political Orientation.

We sampled 20 recent comments from each commenter’s history, for a total of 2000 comments. We then traced the position expressed in the comments. A commenter was considered to show consistency when the political position expressed in all comment samples is consistent. For example, if a commenter expressed her political position in 18 comments out of the 20 samples, the positions of the 18 comments all had to be consistent. Those who changed their position for at least once were tagged as “vague”.

Most of these active commenters showed a consistent political preference, either as liberal or conservative. In the General Set, 92% of the commenters demonstrated consistently either conservative or liberal views; 54% were liberal and 38% were conservative. Similarly, 86% of the commenters in the Popular Set were consistent in their political orientation; 24% were liberal and 62%, conservative. (Table 2)

Regularity of Sentiments

We explore whether commenters showing clear political preferences express different sentiments depending on the political orientation of news articles. If some regularity is observed in their commenting behaviors, it could be potentially used to predict articles’ political orientation.

More specifically, two types of regular behaviors can be considered in the view of prediction regardless of the commenter’s political orientation: leaving a negative comment to the articles revealing a political view opposite to the commenter’s, or a positive comment to those supporting the commenter’s view. For example, a liberal commenter might leave a negative comment on a conservative article and a positive comment on a liberal article. Thus, we define two

types of relationships between the comments’ sentiment and the articles’ political orientation, i.e., Positive and Negative match. Positive match is the case in which the sentiment of a comment is positive and the article’s political orientation is the same with that of the commenter’s. Negative match characterizes a relationship in which the sentiment of a comment is negative and the article’s political orientation is the opposite from that of the commenter.

We observe the degree of regularity in commenter’s behavior according to these two types of relationships. A commenter can be considered highly predictive if he or she frequently shows Positive or Negative match in his or her commenting history. Specifically, a negative comment from a predictive commenter would generate a Negative match with a high probability and a positive comment, Positive match. On the other hand, the commenters who are not predictive would frequently show other relationships rather than Positive or Negative match in their comment history.

This analysis was conducted on the 89 commenters who showed consistent political orientation. Note that consistent political orientation of commenters does not guarantee the regularity of commenting behavior. Although a commenter may have a strong political preference, he or she may express positive or negative sentiments, regardless of articles’ political orientation. For example, a conservative commenter might always blame liberals in his or her comments regardless of articles’ political orientation. It would be difficult to identify articles’ political orientation utilizing such commenters and their comments despite they show consistent political preferences.

We analyzed a part of the commenters’ history to observe the regularity of commenting behavior. We sampled 20 recent comment-article pairs from each of the 89 commenter. For each comment, we manually analyzed the sentiment expressed in the comment, and the political orientation of the article on which the comment was made.

The sentiments of comments were identified as positive, negative, or vague. Examples of positive comments include complimenting or endorsing a politician, showing optimism towards a policy, emphasizing benefits of a policy, etc. The negative types of comments include mockery, criticism and expression of profanity directed towards a party, politician, or a policy. The comments that did not explicitly express positive or negative sentiments were classified as vague. Examples of these comments include proposal of an alternative policy idea, unrelated stories such as advertisements, etc.

The political orientation of articles was identified as liberal, conservative, or vague. Two types of articles were classified as liberal: first, the articles that cover only liberal positions; second, the articles that cover information detrimental to the conservatives. The conservative articles also had two types, vice versa. Other articles were tagged as vague. Examples of these articles include those covering both liberal and

conservative positions, describing general background or implications of an issue, etc.

We analyzed the reliability of both our annotations of comments and news articles. We recruited an annotator who were not aware of our research, and asked to perform the same annotation task. Then, we measured the inter-rater agreement using the kappa measure. The recruited annotator performed annotation on 100 news articles and 100 comments, which were randomly sampled from those we performed annotation on. The kappa measure was 0.73 for the annotations of the news articles and 0.67 for that of the comments. The kappa measure near 0.7 indicates a substantial level of agreement, and the value can be achieved, for example, when 4 out of 5 items are annotated equally [11].

The degree of regularity is calculated for each commenter. This is done by computing two conditional probabilities, the probability of Positive match given a positive sentiment and that of Negative match upon a comment with negative sentiment. Table 3 shows the distribution of commenters according to their conditional probabilities.

	Distribution according to P(T=Negative match S = Negative)	Distribution according to P(T=Positive match S = Positive)
[0.75,1]	18 (G: 6, P: 12)	44 (G: 18, P: 26)
[0.5, 0.75)	43 (G: 16, P: 27)	3 (G: 3, P: 0)
[0.25, 0.5)	24 (G: 20, P: 4)	3 (G: 3, P: 0)
[0, 0.25)	4 (G: 4, P: 0)	2 (G: 2, P: 0)
No sample*	0 (G: 0, P: 0)	37 (G: 20, P: 17)
Total	89 (G: 46, P: 43)	89 (G: 46, P: 43)

G, P: Commenters of the General Set (G) and Popular Set (P)
*: the number of commenters who did not express such sentiment in the samples

Table 3. Distribution of commenters according to $P(T|S)$

T: Type of relations (Negative match, Positive match)
S: Sentiment of a comment (Negative, Positive)

A number of commenters showed highly predictive behaviors that are useful for political view identification. 18 people showed a high level of regularity for the Negative match. When their comment was negative, the Negative match was observed with more than 75% probability. If these commenters leave a negative comment, it is possible to predict that the article’s political orientation is the opposite of the commenters’. 44 people showed regularity for the Positive match.

Behavioral Variability in Commenting

We additionally observed the differences in commenting behaviors among different commenters. Based on the observation, we characterize commenters into three types: *predictive*, *cross*, and *opaque*. As described, the predictive commenters show a high level of regularity.

The cross commenters show less regular behavior than the predictive commenters. This is because they occasionally leave negative comments on the articles showing the same political orientation as their political orientation. Specifically, they leave negative comments on the articles that cast their opponents in a negative light. As these articles are negative toward the opponents, these commenters attack the opponents together in their comments. For example, a

conservative cross commenter often censures Obama in response to the articles criticizing him.

The opaque commenters show low regularity in our analysis. These commenters leave vague comments much more frequently than predictive and cross commenters. In addition, these commenters’ comments are often unrelated to the articles’ topic; for example, opaque commenters talk about their thoughts on other political topics, or even discuss their personal stories. Some commenters wrote and repeatedly copied one comment to many news articles.

SENTIMENT PATTERN-BASED METHODS

In this section, we present sentiment pattern-based methods for political view identification of news stories. Given each commenter, the methods first learn his or her sentiment pattern, i.e., which sentiment he or she expressed depending on the political orientation of news stories. Then, the methods apply it to predict political orientation of the articles on which he or she comments in the future. If the commenter leaves a comment, the methods analyze the expressed sentiment, and predict the political orientation of the article depending on the sentiment pattern.

The key idea behind the methods is to utilize the regularity and consistency in the commenters’ behaviors. From some commenters, we observed a certain level of regularity between the commenters’ sentiment and the political orientation of articles. We also observed that their political preference is consistent over diverse issues.

As mentioned, our methods do not require complex analysis of news text itself. Thus, the methods do not require advanced language analysis tools or rich semantic knowledge of politics.

We developed two methods, the single commenter-based prediction and the multi commenter-based prediction. The single commenter-based prediction utilizes only one commenter’s sentiment pattern for political view identification. On the other hand, the multi commenter-based prediction combines the sentiment pattern of multiple commenters who comment on the same article.

Single Commenter-based Prediction

This method models an individual commenter as a multi-class Bayes’ classifier. It determines the class C of an article as Liberal, Conservative, or Vague, given the sentiment of a comment S , which can be Positive, Negative, or Vague.

$$P(C|S) = \frac{P(C)P(S|C)}{P(S)}$$

$$C = \{Conservative, Liberal, Vague\}$$

$$S = \{Positive, Negative, Vague\}$$

The parameters, i.e., the prior $P(C)$ and the class likelihood $P(S|C)$, are trained from a sample of the comment history of the corresponding commenter. Each training instance is composed of a pair, i.e., comment’s sentiment, and political orientation of the article commented upon.

		Political Orientation (C)		
		Conservative	Liberal	Vague
Sentiment of a comment (S)	Positive	0	0.15	0
	Negative	0.35	0.25	0.2
	Vague	0	0	0.05

Table 4. Example Distribution of $P(C|S)$.

The classifier determines the article’s class as the one which shows the highest posterior probability, i.e., given the sentiment s , it chooses C_i if $P(C_i|S=s) = \max_j P(C_j|S=s)$. Table 4 shows an example of posterior probability $P(C|S)$ distribution. Upon a comment with negative sentiment, the classifier determines the political orientation of the corresponding article as Conservative. Likewise, an article with a positive comment is classified as Liberal.

Multi Commenter-based Prediction

This method utilizes multiple commenters who comment on the same article. It is built upon the single commenter-based method; it aggregates the identification results made from individual commenters. The class of the article is determined based on the aggregated results.

The multi commenter-based prediction is potentially more resistant to prediction errors. The single commenter-based prediction completely depends on a single individual commenter. Thus, some irregular behavior of the modeled commenter may cause prediction error. Through aggregation, the effect of such irregular behavior of a particular commenter can be reduced, and it is more likely that a correct result is achieved. The aggregation of identification results can be done in many ways. We test the two aggregation policies below.

- **Maximum Votes (MV):** Through the single commenter-based prediction, each commenter decides an article’s class as Liberal, Conservative or Other. Maximum Votes aggregates the number of decisions made for each class. The class which obtains the maximum number of decisions is chosen.

choose class i if $V_i = \max(V_{cons}, V_{lib}, V_{vague})$

$$\text{where } V_{cons} = \sum_{c \in C, D(c)=cons} 1,$$

$$V_{lib} = \sum_{c \in C, D(c)=lib} 1, \quad V_{other} = \sum_{c \in C, D(c)=vague} 1$$

C : Commenters of the article $D(c)$: Decision of commenter c

- **Maximum Posterior Probability (MPP):** MPP aggregates the probabilities that the single commenter-based prediction would select each class as the identification results. That is, instead of counting the identification results themselves, it sums up the posterior probability $P(C|S)$ for each class C that is used to make decisions in the single comment-based predictions.

choose class i if $SP_i = \max(SP_{cons}, SP_{lib}, SP_{vague})$

$$\text{where } SP_{cons} = \sum_{c \in C} P_c(C = cons. | S = s_c),$$

$$SP_{lib} = \sum_{c \in C} P_c(C = lib. | S = s_c),$$

$$SP_{vague} = \sum_{c \in C} P_c(C = vague | S = s_c)$$

s_c : The sentiment expressed by the commenter c .

$P_c(C|S)$: Posterior probability calculated from commenter c ’s sentiment pattern.

Simple Sentiment Classifier

The proposed methods take the sentiment of comments as their input. Thus, a sentiment classifier is required to automate the political view identification process. We developed a simple sentiment classifier.

The simple sentiment classifier adopts the language model approach [23]. The approach requires a training set of comments for each class. From the training set, it constructs a probabilistic model of how words are likely to appear according to their class. The probability that a term w would appear from a class c_i is calculated as:

$$P(w|c_i) = \frac{tf(w, c_i)}{|c_i|}$$

where $tf(w, c_i)$ represents the frequency of word w in the training set of comments for class c_i , and $|c_i|$ is the total number of terms in the training set for class c_i .

The simple sentiment classifier determines the class based on the developed model. It selects the class that produces the largest probability value. The probability that a comment r_k belongs to c_i is the product of the probability $P(w|c_i)$ of its words. The comment may include words that are unseen in the training set. For such words, we assigned a non-zero probability following the Laplace smoothing technique [14].

$$P(r_k|c_i) = \prod_{w \in r_k} P(w|c_i)$$

We have trained and evaluated the simple sentiment classifier with the comments of the commenters selected in our analysis. The 5-fold cross validation showed 73.6% accuracy.

EVALUATION

We evaluated the effectiveness of the proposed methods for political view identification. For the evaluation, the methods used the commenters who showed consistent political orientation in our analysis; 46 commenters sampled from the General Set, and 43 commenters sampled from the Popular Set. The evaluation is twofold: first, the accuracy of the prediction; second, the article coverage of the methods. The accuracy is the proportion of correct predictions on the article’s political orientation among all predictions. The article coverage is the proportion of articles commented upon by these commenters.

Comparison Method

We compare the performance of the proposed method to a news text analysis-based method. For the comparison method, we adopt a commonly used approach for text classification, TF/IDF for feature extraction and SVM algorithm for classification.

The comparison method takes the words of articles as features. In order to obtain quality features, we removed stop words (e.g., pronouns, conjunctions) and auxiliary words

(e.g., postposition) through a morphological analysis [26], and removed irrelevant contents such as advertisements. The weight of each word is calculated by following the TF/IDF weighting scheme [20]. For classification, we used the implementation of the SVM multiclass classifier [27]. It predicts the class of an article as Liberal, Conservative, or Vague. The classifier used the linear kernel and a regularization parameter was optimized through 5-fold cross validation.

Evaluation of Single Commenter-based Prediction

The Bayes classifier for each commenter is trained using 20 comment-article pairs. For accuracy evaluation, we also developed a test set for each classifier. The test set includes the most recent 20 comment-article pairs of the corresponding commenter.

The benefits of single commenter-based prediction can be observed from the commenters who achieve high accuracy. Thus, we first focused our evaluation to the commenters achieving accuracy over 70%. In both the General Set and the Popular Set, 7 commenters satisfy the condition (Table 5 shows the distribution of the commenters according to the accuracy of the classifiers). We denote them as predictive commenters (PC).

	# of Commenters
[0.7,1]	14 (G: 7, P: 7)
[0.6, 0.7)	25 (G: 14, P: 11)
[0, 0.6)	50 (G:25, P: 25)
Total	89 (G: 46, P: 43)

Table 5. Distribution of Commenters According to Accuracy.

We evaluate two versions of the proposed method. In the first version, denoted as MA, the sentiment of comments is manually analyzed. This version shows the accuracy upper bound of the method. The second version, denoted as SA, uses the simple sentiment classifier to automatically analyze the sentiment of comments. The accuracy of these methods is compared to the news text analysis-based method, denoted as TA. The TA method is trained and tested with the same training set and test set.

Table 6 shows the accuracy of the predictions on articles' political orientation. In both General Set and Popular Set, the proposed method outperforms the TA method. The accuracy of SA is lower than MA as the simple sentiment classifier misidentifies the expressed sentiment for some comments. Despite the limitations of the simple sentiment classifier, SA achieves better performance than TA.

	General Set			Popular Set		
	MA	SA	TA	MA	SA	TA
Overall	76%	67%	48%	74%	66%	45%
Conservative or Liberal	83%	75%	56%	80%	75%	51%
Vague	13%	0%	8%	25%	0%	0%

Table 6. Accuracy of predictive commenters.

The accuracy of the proposed method is higher when considering only the Conservative and Liberal predictions.

This is because the predictive commenters show a more regular commenting behavior towards Conservative and Liberal articles than towards Vague articles. Another reason is that there are more Conservative and Liberal articles in the training set than Vague articles. Near 90% of the articles in the training set were either Conservative or Liberal. Thus, it is easier to observe patterns for Conservative or Liberal articles from the training set. As the commenters have a clear political preference, they respond more actively to Conservative and Liberal articles than to Vague articles. Although the accuracy of the method is poor for predicting Vague articles, this does not affect the overall accuracy significantly. We believe many applications will have more interest in identifying Conservative and Liberal articles than Vague articles. Such applications can exploit the proposed method by taking only the predictions made as Conservative or Liberal.

	General Set	General Set (≥ 5 comments)	Popular Set
PC	5%	21 %	31%
PC+CC	12%	46 %	52%

Table 7. Article Coverage of the Selected Commenters.

We present the article coverage of the predictive commenters in the second row of Table 7. Although the method uses a small number of commenters, it covers a considerable number of articles. The commenters of the Popular Set cover nearly 40% of the articles of the set. Those of the General Set cover fewer articles; however, this is because the set has a small number of articles with many comments. When considering only the articles with more than 5 comments, the article coverage reaches 21%.

	General Set			Popular Set		
	MA	SA	TA	MA	SA	TA
Overall	67%	56%	42%	67%	57%	48%
Conservative or Liberal	75%	69%	51%	74%	68%	57%
Vague	37%	5%	12%	30%	0%	0%

Table 8. Accuracy after adding cross commenters.

We evaluate the proposed method after including more commenters. We additionally use the commenters whose accuracy scores between 60% and 70%. Most of them were cross commenters (CC). 14 more commenters from the General Set, and 11 more from the Popular Set are used. 39 Bayes classifiers are developed in total. Table 8 shows the accuracy achieved after adding the cross commenters. The accuracy drops about 8% as the added commenters show lower accuracy. Although the accuracy drops, the proposed method still performs better than the TA method.

Adding the cross commenters enlarges the article coverage (refer to the third row of Table 7). More than half of the articles are covered in the Popular Set. The article coverage increases more than twice for the General Set. When considering only the articles with more than 5 comments, 46% of the articles are covered.




Figure 6. Accuracy Depending on Training Set Size.

We also observed the effect of the training set size on accuracy. We measured the overall accuracy while varying the training set size from 5 comment-article pairs to 20 comment-article pairs. As shown in Figure 6, the training set size does not cause much variation in the accuracy. This indicates that the proposed method does not require many training data.

Evaluation of Multi Commenter-based Prediction

The multi commenter-based prediction is only evaluated for the articles of the Popular Set (3385 articles). Recall that the method aggregates the results of multiple commenters. Thus, the benefits of the method can be observed when many commenters comment on the same articles. However, such articles are rarely found in the General Set.

The method makes prediction by aggregating the results of the 43 commenters of the Popular Set. We measured the accuracy while varying the minimum number of commenters required for the prediction, from 1 to 12. When the number is set to 3, for instance, the method makes prediction for the articles upon which at least three commenter out of 43 commented. In this evaluation, the sentiment of comments was analyzed only with the simple sentiment classifier. Due to the large scale of comments (8871 comments), we could not manually identify the sentiment of the comments.

We measured the accuracy for the two aggregation policies, Maximum Votes, and Maximum Posterior Probability. We compared the accuracy of these two policies to that of the TA method (See Figure 7).




Figure 7. Accuracy of Multi Commenter-based Prediction (Overall)

The two aggregation policies perform better as more commenters' predictions are aggregated. For the articles having more than 12 commenters' comments, the two methods all achieves near 70% accuracy. On the other hand, the accuracy of the TA method is constantly lower than 50%. The accuracy of the proposed methods increases when considering only the Conservative and Liberal predictions

(refer to Figure 8). For the articles with more than 12 commenters' comments, all two aggregation methods achieved accuracy over 80%.




Figure 8. Accuracy of Multi Commenter-based Prediction (Conservative and Liberal Predictions)

Figure 10 shows the article coverage depending on the minimum number of commenters required for prediction. There is a tradeoff between accuracy and article coverage. The coverage is large when predictions are made even with small number of commenters. However, the accuracy is not high and similar to that of the TA method. The article coverage decreases as the method require more number of commenters, however, the accuracy improves.




Figure 10. Article Coverage of Multi Commenter-based Prediction

DISCUSSION AND FUTURE WORK

In this paper, we discussed the effectiveness of commenters' sentiment pattern for political view identification of news articles. The performance of the proposed method seems promising to be used in real world scenarios. The method can be practically used for applications which can tolerate a few errors, e.g., news recommendation, search; for example, a list of liberal or conservative news items can be made for recommendation with near 80% precision through predictive commenters. It is also important that the method supports applications to selectively use good results, i.e., predictions made with high probability, e.g., the predictions made through multiple predictive commenters.

Applying the idea for identifying more nuanced political positions is also possible. More detailed characterization of the positions, such as issue, politician, or region-based positions, and searching for commenters taking such positions would be helpful. More advanced modeling of commenters' behavior is another direction, e.g., modeling collective behavior of liberals and conservatives may help detecting another position with which both groups disagree.

In addition, we believe similar approach can be made for other discourse domains where many people share and express different preferences, such as celebrity news, product reviews. There are likely to be predictive commenters who react negatively to stories on certain artists or products but positively to others.

There is also room for further improvement of the proposed method in terms of accuracy and coverage. A potential approach for improvement is to discover more predictive commenters. Our future work includes investigating solutions for discovering predictive commenters. The discovering process involves three stages: discover politically biased commenters, check if they express explicit sentiments, and analyze their sentiment patterns. Approaches can be made for each stage. Targeting the search for biased commenters, e.g., to popular biased blogs, would reduce the search space. Automatic tools can be used to check explicit sentiments in their comments. Analyzing the similarity of their sentiment patterns to already-discovered predictive commenters would be possible.

It would be interesting to discover more useful features from other social annotations and combine them. Advanced methods for sentiment analysis [17] can be also adopted.

REFERENCES

1. Adamic, L. A. et al., The political blogosphere and the 2004 U.S. election: divided they blog. In *Proc. of the 3rd international workshop on Link discovery*, 2005.
2. Baker, B. How to Identify, Expose and Correct Liberal Media Bias. Media Research Center, 1994.
3. Bernhardt, D., et al. Political Polarization and the Electoral Effects of Media Bias. CESifo Working Paper Series No. 1798.
4. Della Vigna, S. and Kaplan, E. The Fox News effect: Media bias and voting. University of California, Berkeley Mimeograph. 2005.
5. Entman R.M. Framing: Toward Clarification of a Fractured Paradigm. *Journal of Communication* 43 (4): 51-8, 1993.
6. Gamon, M., et al., Blews: Using blogs to provide context for news articles, In *Proc. of ICWSM*, 2008.
7. Gilbert, E., et al. Blogs Are Echo Chambers: Blogs Are Echo Chambers, In *Proc. of HICSS*, 2009.
8. Jeong, I., Analyzing Editorials of 10 News Providers: Clear Antagonistic Relationship between Liberal vs. Conservative, and Newspaper vs. Broadcast. *Newspaper and Broadcasting*, Korea Press Foundation, 2008.
9. Jiang, M. et al., Exploiting subjectivity analysis in blogs to improve political leaning categorization, In *Proc. of ACM SIGIR*, 2008.
10. Kull, S., C. Ramsey and E. Lewis. Misperceptions, the media, and the Iraq War. *Political Science Quarterly* 118(4): 569-598, 2003.
11. Landis JR, Koch G. The measurement of observer agreement for categorical data. *Biometrics* 33:159-174. 1977.
12. Lin, W.H., Identifying Ideological Perspectives in Text and Video, PhD thesis, Language Tech. Inst., School of Computer Science, Carnegie Mellon University, 2008.
13. G. Mishne and N. Glance, Leave a reply: An analysis of weblog comments, In *Third Annual Workshop on the Weblogging Ecosystem: Aggregation, Analysis and Dynamics*, 2006.
14. Mitchell, T.M., Machine Learning, McGraw Hill. 1997.
15. Munson, S. et al. Sidelines: An Algorithm for Increasing Diversity in News and Opinion Aggregators. In *Proc. of ICWSM*, 2009.
16. Oh, A. et al. User Evaluation of a System for Classifying and Displaying Political Viewpoints of Weblogs. In *Proc. of ICWSM*, 2009.
17. Pang, B., et al., Opinion Mining and Sentiment Analysis. *Foundations and Trends in Information Retrieval* 2(1-2), pp. 1-135, 2008.
18. Park, S. et al. NewsCube: Delivering Multiple Aspects of News to Mitigate Media Bias. In *Proc. of CHI*, 2009.
19. Park, S. et al. Aspect-level News Browsing: Understanding News Events from Multiple Viewpoints. In *Proc. of IUI*, 2010.
20. Salton, G. and M. J. McGill. Introduction to modern information retrieval. McGraw-Hill. 1983.
21. Tankard, James W. The Empirical Approach to the Study of Media Framing. In *Framing Public Life: Perspectives on Media and our Understanding of the Social World*. Stephen D Reese, Oscar H Gandy, and August E Grant. Mahwah, NJ: Lawrence Erlbaum Associates. 2001.
22. Tsagkias, M., et al., News Comments: Exploring, Modeling, and Online Prediction, In *Proc. ECIR*, 2010.
23. Yamron, JP., et al., Statistical models of topical content, *The Kluwer International Series On Information Retrieval*, pages 115-134, 2002.
24. Yano, T., et al., What's Worthy of Comment? Content and Comment Volume in Political Blogs", In *Proc. of ICWSM*, 2010.
25. Yee, W. G., et al., Are Web User Comments Useful for Search?, In *Proc. of LSDS-IR at SIGIR*, 2009.
26. Korean Morphological Analyzer MACH 1.0. Available online at <http://cs.sungshin.ac.kr/~shim/demo/mach.html>
27. SVM multiclass classifier. Available online at http://svmlight.joachims.org/svm_multiclass.html